

Fast and Reliable Primary Frequency Reserves From Refrigerators with Decentralized Stochastic Control

Evangelos Vrettos^{*1}, Charalampos Ziras^{†2}, and Göran Andersson^{‡1}

¹Power Systems Laboratory, ETH Zurich, Switzerland

²Department of Electrical Engineering, Technical University of Denmark
(DTU), Lyngby, Denmark

October 5, 2016

Abstract

Due to increasing shares of renewable energy sources, more frequency reserves are required to maintain power system stability. In this paper, we present a decentralized control scheme that allows a large aggregation of refrigerators to provide Primary Frequency Control (PFC) reserves to the grid based on local frequency measurements and without communication.

The control is based on stochastic switching of refrigerators depending on the frequency deviation. We develop methods to account for typical lock-out constraints of compressors and increased power consumption during the startup phase. In addition, we propose a procedure to dynamically reset the thermostat temperature limits in order to provide reliable PFC reserves, as well as a corrective temperature feedback loop to build robustness to biased frequency deviations. Furthermore, we introduce an additional randomization layer in the controller to account for thermostat resolution limitations, and finally, we modify the control design to account for refrigerator door openings.

Extensive simulations with actual frequency signal data and with different aggregation sizes, load characteristics, and control parameters, demonstrate that the proposed controller outperforms a relevant state-of-the-art controller.

^{*}vrettos@eeh.ee.ethz.ch

[†]chazi@elektro.dtu.dk

[‡]andersson@eeh.ee.ethz.ch

CDF Cumulative Distribution Function
COP Coefficient of Performance
PDF Probability Distribution Function
PFC Primary Frequency Control
RES Renewable Energy Sources
RMSE Root Mean Square Error
MAPE Mean Absolute Percentage Error
MHSE Moving Horizon State Estimator
TCL Thermostatically Controlled Load

Table 1: Nomenclature: symbols

Symbol	Unit	Description
C	$\text{kJ}/^\circ\text{C}$	Refrigerator thermal capacitance
e_r	%	Instantaneous reserve error
$e_{b,\text{mape}}$	%	Baseline MAPE
$e_{r,\text{mape}}$	%	Reserve MAPE
$e_{p,\text{rmse}}$	-	Switching probability RMSE
D	-	Duty cycle
D^a	-	Actual duty cycle (with PFC)
D^n	-	Nominal duty cycle (without PFC)
D^d	-	Desired duty cycle (with PFC)
D^r	-	PFC reserve capacity in terms of duty cycle
E_{cl}^d	kJ	Daily energy consumption without door openings
E_{op}^d	kJ	Daily energy consumption with door openings
F_{off}	-	CDF of the lock-off time in the aggregation
F_{on}	-	CDF of the lock-on time in the aggregation
K_c	-	Corrective temperature gain
$K_{r,t}$	$^\circ\text{C}$	Resetting factor for the thermostat limits
l_{off}	-	Lock-off event as a random variable
l_{on}	-	Lock-on event as a random variable
$L_{\text{off}}^{\text{st}}$	-	Steady-state fraction of refrigerators locked-off
$L_{\text{on}}^{\text{st}}$	-	Steady-state fraction of refrigerators locked-on
$L_{\text{off},t}^{\text{tr}}$	-	Transient fraction of refrigerators locked-off
$L_{\text{on},t}^{\text{tr}}$	-	Transient fraction of refrigerators locked-on
m	-	Refrigerator on/off state

1 Introduction

1.1 Motivation and Related Work

It is expected that the increasing shares of intermittent Renewable Energy Sources (RES) will increase the need for ancillary services in power systems [1]. Since RES displace conventional power plants, there is a growing interest in exploiting the flexibility of demand-side resources to provide ancillary services, namely frequency and voltage control [2]. Thermostatically Controlled Loads (TCLs) with thermal inertia such as space and water heaters, air conditioners and refrigerators are well suited for provision of ancillary services, because their consumption can be shifted in time without user discomfort.

The idea of controlling loads to provide ancillary services is not new. In 1980, [3] proposed a simple frequency-responsive controller to adjust the temperature limits of TCLs based on frequency measurements to provide a governor-type action. Much of the early work focused on modeling populations of TCLs based on first principles [4–6]. For example, [5] proposed a well-known model consisting of a set of Fokker-Planck equations that describe the Probability Distribution Function (PDF) of temperature in a TCL aggregation.

The first line of research on load control has focused on centralized control of TCL aggregations to provide load following or frequency regulation. References [7–9] developed controllers to provide such services via thermostat setpoint

Table 2: Nomenclature: symbols (continued)

Symbol	Unit	Description
N_{ev}	s	Duration of a step frequency deviation event
N_{on}	s	Maximum lock-on time within the aggregation
N_{off}	s	Maximum lock-off time within the aggregation
N_r	-	Size of refrigerator aggregation
N_{rec}	s	Recovery period after a step frequency deviation
$N_{r,a}$	-	Number of refrigerators activated for PFC
N_s	s	Duration of startup dynamics
N_{sim}	s	Simulation period
N_δ	s	Half-period of the frequency deviation signal
P_{agg}	W	Actual aggregate electric power
P_b	W	Uncontrolled aggregate electric power
P_d	W	Desired aggregate electric power (with PFC)
P_{res}	W	Reserve capacity of the aggregation
P_n	W	Nominal refrigerator electric power
q	-	Switching probability of the benchmark controller
R	$^{\circ}\text{C}/\text{kW}$	Refrigerator's thermal resistance with closed door
R_{op}	$^{\circ}\text{C}/\text{kW}$	Refrigerator's thermal resistance with open door
r	-	Normalization factor for the resetting factor
S_{off}	-	Survival function of the lock-off time in the aggregation
S_{on}	-	Survival function of the lock-on time in the aggregation
S_u	-	Refrigerator startup dynamics profile
$s_{net,t}$	-	Average net switching rate in the aggregation
$s_{off,t}$	-	Average rate at which refrigerators switch off
$s_{on,t}$	-	Average rate at which refrigerators switch on

manipulation, and showed that linear models are sufficient to describe the aggregate transient dynamics of TCLs under thermostat setpoint control.

Other researchers developed control strategies based on probabilistic switching, where a fraction of the devices switches stochastically to provide the desired service. Reference [10] showed how bi-directional control can be achieved via broadcasting a signal that only switches off devices. Broadcast controllers were proposed in [11] and [12] to allow the aggregate power of a population of air conditioners and refrigerators to track a power reference signal. Reference [13] proposed a hierarchical load control framework, where the upper layer is centralized and computes optimal control gains for loads at different buses, and the lower layer is decentralized with switching probabilities that depend on the control gains.

The above works as well as other relevant approaches that rely either on probabilistic switching or rule-based controllers (e.g., [14–16]) assume access to TCL state measurements (temperatures and on/off states). More recent works have proposed state estimation methods to reduce the needs for real-time communication. For example, [17] used a Kalman Filter to estimate the state of a Markov model for a TCL aggregation, whereas [18] used a similar modeling approach and a Moving Horizon State Estimator (MHSE). Reference [19] developed a particle filter, whereas [20] proposed a state estimator based on partial differential equations. The authors of [21] developed an MHSE to reconstruct the states of individual TCLs using aggregate power measurements from substations and low-frequency

Table 3: Nomenclature: symbols (continued)

Symbol	Unit	Description
T	$^{\circ}\text{C}$	Refrigerator temperature
T_a	$^{\circ}\text{C}$	Ambient temperature
T_{\max}	$^{\circ}\text{C}$	Higher deadband limit of thermostat
T_{\min}	$^{\circ}\text{C}$	Lower deadband limit of thermostat
\bar{T}_{nom}	$^{\circ}\text{C}$	Nominal mean temperature (without PFC)
T_{set}	$^{\circ}\text{C}$	Thermostat setpoint [$T_{\text{set}} = 0.5 \cdot (T_{\max} + T_{\min})$]
\dot{T}_d	$^{\circ}\text{C/s}$	Average temperature decrease rate
\dot{T}_i	$^{\circ}\text{C/s}$	Average temperature increase rate
\bar{T}_{cl}	$^{\circ}\text{C}$	Average refrigerator temperature with closed door
\bar{T}_{op}	$^{\circ}\text{C}$	Average refrigerator temperature with open door
t_{off}	s	Duration of refrigerator's off cycle
t_{off}^l	s	Refrigerator lock-off time
t_{on}	s	Duration of refrigerator's on cycle
t_{on}^l	s	Refrigerator lock-on time
u	-	Peak power factor for startup dynamics modeling
w	$^{\circ}\text{C/s}$	Noise term for refrigerator's external disturbances
x	-	Fraction of switched loads for PFC under startup dynamics
z	-	Thermostat resetting event as a Bernoulli random variable

TCL state measurements from smart meters.

The second line of research investigated how TCLs can provide frequency support in a decentralized way based on local frequency measurements. References [22] and [23] developed deterministic approaches that rely on a frequency-dependent temperature deadband to provide Primary Frequency Control (PFC). Reference [24] introduced a more generic approach that is based not only on frequency deviation measurements, but also on their evolution over time. A simple rule-based controller with delays was used in [25] in an experimental demonstration to provide PFC with residential appliances.

Such deterministic approaches provide an effective initial response and reduce the frequency excursions after a sudden disturbance. Despite their simplicity, these approaches have an important limitation: they cause rebound effects and tend to synchronize the on/off cycles of individual devices, which might introduce non-decreasing oscillations in frequency [26, 27].

In order to de-synchronize the responses of TCLs while providing PFC, a few randomized approaches have been proposed in the literature. For example, a control approach with random on/off frequency thresholds and turn-on delay times was used in an experimental demonstration in [28]. Reference [29] proposed a randomized load control scheme where each load monitors the frequency over random time instances and responds according to a simple control policy based on frequency thresholds.

Reference [26] proposed an interesting stochastic approach to control a population of refrigerators. Each device is modeled as a Markov-jump linear system with transition probabilities between the on and off states that depend indirectly on the frequency deviation. The authors derived closed form expressions for the mean value and the variance of the average temperature within the population, and

Table 4: Nomenclature: Greek letters

Symbol	Unit	Description
α	1/s	Thermal parameter of continuous-time model
β	$^{\circ}\text{C}/\text{kJ}$	Thermal parameter of continuous-time model
γ	$^{\circ}\text{C}/\text{Hz}$	Auxiliary variable for PFC
ΔD	-	Duty cycle change
ΔD_t^{op}	-	Increase in duty cycle due to door openings
ζ	-	Auxiliary variable for thermostat resetting events
Δf	Hz	Frequency deviation
ΔT	$^{\circ}\text{C}$	Width of thermostat's deadband
ΔT_b	$^{\circ}\text{C}$	Bound for refrigerator's thermostat limits deviation
ΔT_{lim}	$^{\circ}\text{C}$	Change in refrigerator's thermostat limits
ΔT_{res}	$^{\circ}\text{C}$	Minimum resolution of refrigerator's thermostat
Δt	s	Discretization time step
δ	Hz	Magnitude of frequency deviation signal
ε	$^{\circ}\text{C}$	Tolerance for temperature deviation during a frequency event
ϵ	$^{\circ}\text{C}$	Tolerance for temperature deviation after a frequency event
η	-	Coefficient of performance
λ	-	Auxiliary variable equal to $1 - K_c$
μ_d	s	Mean duration of a door opening event
μ_{op}	-	Average number of door openings per day
ν	-	Parameter for modeling of startup dynamics
ξ	-	Normalized energy consumption increase due to door openings
ρ	-	Switching probability of the proposed controller
ϱ	-	Switching probability for resetting of thermostat limits
σ_d	s	Standard deviation of the duration of a door opening event
σ_{op}	-	Standard deviation of the number of door openings per day
χ	-	Switching action of a refrigerator (Bernoulli random variable)

Table 5: Nomenclature: probability operators

Operator	Description
$\mathbb{E}[\cdot]$	Expected value of a random variable
\mathbb{N}	Normal probability distribution
$\mathbb{P}[\cdot]$	Probability of a random variable
$\text{SD}[\cdot]$	Standard deviation of a random variable
\mathbb{U}	Uniform probability distribution
$\text{Var}[\cdot]$	Variance of a random variable

proved that the closed-loop system is asymptotically stable. However, the control adjusts the properties of the steady state temperature distribution, and therefore results in slow responses that are unacceptable for PFC.

Another relevant decentralized stochastic controller was developed in [30]. The main innovation of this work is that the average heating rate of the population is used as a control variable. The formulations of [30] result in control laws for individual appliances, namely temperature limit changes and switching rates, and allow an aggregator to estimate the available flexibility from the aggregation. However, the control approach of [30] is able to respond to demand reduction requests but not demand increase requests, and thus it is not appropriate for PFC.

More related to our work is [31] where a decentralized stochastic controller based on probabilistic switching is proposed for refrigerators to provide PFC. This

Table 6: Nomenclature: subscripts and superscripts

Subscript/Superscript	Description
i	Refrigerator index within an aggregation
min/max	Minimum/maximum value of a variable or parameter
k, t	Time indexes
$(\bar{\cdot})$	Average value of a variable or parameter
$(\hat{\cdot})$	Estimated value of a variable or parameter
$[\cdot]_+$	Caps a variable to non-negative values

controller results in fast responses that are suitable for PFC, but it has a number of limitations that will be discussed in Section 1.2.

Some papers, for example [32–35], proposed methods to provide PFC at minimum cost or maximum social welfare. Note that we do not address this topic and, therefore, the proposed controller is complementary to the work presented in [32–35]. Reference [32] proposed an optimization problem formulation to allocate trigger frequencies to on/off devices such that they collectively provide PFC reserves at minimum cost. However, [32] used a simplified model for on/off devices, which does not model TCLs in full detail. Reference [33] proposed a frequency support method for residential loads based on utility functions to maximize social welfare. The authors of [34,35] formulated a load control optimization problem and developed synchronous and asynchronous algorithms to solve it in a decentralized way. References [34,35] used simplified aggregate load models for each bus of the transmission system, but did not propose methods to coordinate a population of on/off devices.

1.2 Contribution and Organization of this Paper

In this paper, we develop a new decentralized stochastic method to provide accurate and reliable PFC reserves from an aggregation of refrigerators.¹ In contrast to the centralized approaches of [10–14, 16], the proposed controller relies on switching probabilities that are computed based on local frequency measurements, and therefore there is no need for real-time communication. Our work is similar to [26, 29, 30] in that it randomizes the reaction of individual loads to avoid synchronization. The main advantage of probabilistic switching is a much faster PFC response compared with [29] where the time instance of the reaction of each load is randomized, and [26, 30] where the devices’ switching rates depend on the desired evolution of mean temperature.

A similar probabilistic switching approach for PFC with refrigerators was proposed in [31]. This approach relies on four main assumptions: (i) each refrigerator

¹Although we consider residential refrigerators, the proposed approach can be applied to other TCLs with compressors and similar dynamics. The main advantage of residential refrigerators over other TCLs is that they continuously cycle on and off and that their power consumption has very little seasonality (due to regulated room temperature). Therefore, the potential for PFC reserve is relatively constant all year round. Refrigerators are more well-suited for PFC than secondary and tertiary control, due to their limited energy capacity.

consumes a constant amount of power at the on state; (ii) each refrigerator can cycle on and off arbitrarily often; (iii) the frequency deviation signal is zero-mean; and (iv) the refrigerators are not exposed to door openings. Assumptions (i), (ii) and (iv) were also made in [30].

In practice, the power consumption of a refrigerator is typically higher at the beginning of each on cycle (the so-called startup dynamics), its compressor has lockout constraints and cannot turn on and off arbitrarily often, and its temperature is affected by door openings. Furthermore, the frequency deviation might be significantly biased to the positive or negative direction for a few hours. Since these assumptions do not hold in practice, the method of [31] might fail to provide accurate PFC reserves in a realistic environment, and thus its applicability is questionable.

Our main contribution is a new controller that works in a realistic environment without the limiting assumptions (i) - (iv).² Apart from the basic probabilistic switching formulation which is similar to that of [31], the proposed controller consists of several novel components that allow us to drop assumptions (i) - (iv). Furthermore, we show that combining the probabilistic switching with deterministic control of thermostat temperature limits improves the control performance. Nevertheless, if the frequency deviation signal is biased, the controller's performance will deteriorate because the mean temperature of the population will deviate significantly from the nominal value. For this purpose, we develop an additional control loop to regulate the mean temperature of the population, and derive analytical upper and lower bounds for the control gain. In addition, we introduce another randomization layer in the controller to compensate for the typically limited thermostat resolution. We analytically show that the side-effect is an increased variance of the thermostat limit changes which, nevertheless, can be bounded. In comparison with the controller of [31], the proposed controller achieves a significant improvement in the continuous operation of a refrigerator aggregation under PFC provision, as well as a small improvement in dynamic response.

The remainder of this paper is organized as follows. Section 2 introduces the refrigerator model and probabilistic switching control approach. In Section 3, we extend the probabilistic switching to account for startup dynamics and lockout constraints. Section 4 discusses the advantages and implications of controlling the thermostat temperature limits, whereas in Section 5 we integrate refrigerator door openings in the control scheme. In Section 6, we introduce the model parameters and performance metrics, which are used in the simulations and sensitivity analysis of Section 7. In Section 8, we show the controller's robustness to biased frequency deviations, limited thermostat resolution, and door openings. Implementation issues are discussed in Section 9, whereas Section 10 concludes.

²This paper extends, improves and formalizes the methods presented in our preliminary work [36].

2 Modeling and Probabilistic Switching

2.1 Refrigerator Model

Consider a refrigerator i without freezer controlled by an on/off hysteresis controller. We use a standard first order differential equation to model the lumped temperature $T_i(t)$ of the refrigerator including the air and solid mass (refrigerator compartment and content) [6, 26]

$$\dot{T}_i(t) = \alpha_i [T_{a,i} - T_i(t)] + w_i(t) \quad \text{when OFF} \quad (1)$$

$$\dot{T}_i(t) = \alpha_i [T_{a,i} - T_i(t)] - \beta_i P_{n,i} + w_i(t) \quad \text{when ON,} \quad (2)$$

where $\alpha = 1/RC$, $\beta = \eta/C$, C is the thermal capacitance, R is the thermal resistance, T_a is the room temperature, η is the Coefficient of Performance (COP), and P_n is the nominal power. The noise term $w(t)$ aggregates the effect of external disturbances, e.g., door openings, changes in food content, and variations of T_a . In Section 5, we model the effect of door openings that are the dominant disturbances.

The model is discretized with a discretization time step $\Delta t = 1$ sec. Let $T_{\min,i}$, $T_{\max,i}$ denote the deadband limits of the thermostat's hysteresis controller and $m_{i,t} \in \{0, 1\}$ denote the compressor's on/off state. Using the discrete time model we can derive the duration of the on and off cycles (t_{on} and t_{off})

$$t_{\text{on},i} = R_i C_i \ln \left[\frac{T_{\max,i} - T_{a,i} + \eta_i R_i P_{n,i}}{T_{\min,i} - T_{a,i} + \eta_i R_i P_{n,i}} \right] \quad (3)$$

$$t_{\text{off},i} = R_i C_i \ln \left[\frac{T_{\min,i} - T_{a,i}}{T_{\max,i} - T_{a,i}} \right], \quad (4)$$

and the device's duty cycle $D_i = t_{\text{on},i} / (t_{\text{on},i} + t_{\text{off},i})$.

2.2 Aggregation Model

Consider a population of N_r refrigerators providing PFC. For control purposes we are interested in the aggregate power of the population $P_{\text{agg},t}$, which depends on the aggregate duty cycle D_t , i.e., the fraction of loads that are at the on state. If N_r is sufficiently large, D_t and $P_{\text{agg},t}$ are given by

$$D_t = (1/N_r) \cdot \sum_i m_{i,t}, \quad P_{\text{agg},t} \approx N_r D_t \bar{P}_n, \quad (5)$$

where $i \in [1, N_r]$ is the refrigerator index and $\bar{\cdot}$ denotes the mean value of a parameter in the aggregation. Without PFC and if $w_i(t) = 0 \forall i$, the duty cycle will be approximately constant and equal to a nominal value D^n , whereas the aggregate power will be approximately equal to $P_{\text{agg},t} \approx N_r D^n \bar{P}_n$.

2.3 Primary Frequency Control with Probabilistic Switching

PFC reserves can be provided by perturbing the population's duty cycle around D^n proportionally to the frequency deviation Δf_t with a gain D^r , which corresponds to the PFC reserve capacity (the reserve capacity increases with the value of D^r). Therefore, in order to respond to a frequency deviation Δf_t the duty cycle should be controlled to the desired value

$$D_t^d = D^n + D^r \cdot \frac{\Delta f_t}{\Delta f_{\max}}, \quad (6)$$

where Δf_{\max} is the frequency deviation for full PFC reserve activation, e.g., 0.2 Hz in continental Europe.

Let us denote by $\Delta D_t = D_t^d - D_{t-1}^d$ the change in desired duty cycle between two consecutive time steps. Note that ΔD_t depends on Δf_t and Δf_{t-1} , which can be measured locally by each refrigerator, and so the duty cycle change can be achieved in a decentralized way with probabilistic switching. Each device calculates the switching probability

$$q_t = \begin{cases} \frac{\Delta D_t}{1 - D_{t-1}^d}, & \text{if } \Delta D_t \geq 0 \text{ (switch on)} \\ -\frac{\Delta D_t}{D_{t-1}^d}, & \text{if } \Delta D_t < 0 \text{ (switch off)}, \end{cases} \quad (7)$$

and generates a random number uniformly distributed in $[0, 1]$. If the number is smaller than q_t , the device will switch to contribute to PFC. The first line of (7) is the probability for devices that are off to switch on; thus, the desired duty cycle change ΔD_t is normalized by the portion of devices that are currently off ($1 - D_{t-1}^d$). The second line of (7) is the probability for devices that are on to switch off, and ΔD_t is normalized by the portion of devices that are currently on (D_{t-1}^d).

If each refrigerator responds to the switching probability (7), the fraction of switched devices will be very close to ΔD_t due to the law of large numbers. Therefore, the population will collectively provide the required reserve without any real-time communication between the loads.

2.4 Offline Communication and Coordination

Although no real-time communication is required to provide PFC, the refrigerators need some offline coordination. The main reason for this is to fix the value of D^r that reflects the PFC reserve capacity, which is typically determined in a reserve market on a weekly or daily basis.³ Furthermore, as discussed in the following sections, the control design relies on mean values of several parameters in the population (e.g., \bar{P}_n). Therefore, if the statistical properties of the population change

³A refrigerator aggregation can provide PFC in parallel to other reserve providers (e.g., generators or demand-side resources) without any coordination with them other than the market-based determination of reserve capacities.

over time, the respective controller parameters should be updated. A communication channel with low data transfer can be used to (infrequently) coordinate the refrigerators by communicating the updated controller parameters and D^r values.

3 Startup Dynamics and Lockout Constraints

3.1 Refrigerator Startup Dynamics

Typically, a refrigerator's compressor consumes more power during the startup phase due to a smaller COP, higher evaporator temperature, and higher motor current [37]. Since this additional power is up to 25% of the nominal power P_n , it is important to consider it in PFC to avoid overshoots and large reserve errors.

We model the startup dynamics of refrigerator i using a peak power factor $u_i > 0$ and the duration of the startup phase $N_{s,i}$. We assume that the power becomes immediately equal to $(1 + u_i)P_{n,i}$ when the refrigerator switches on, and then it linearly decreases until it becomes equal to $P_{n,i}$ after $N_{s,i}$ seconds. Thus, the refrigerator power evolves according to

$$P_{i,t} = P_{n,i}[1 + S_{u,i}(t)] = P_{n,i} \left[1 + u_i \left[1 - \frac{t}{N_{s,i}} \right]_+ \right], \quad (8)$$

where $S_{u,i}$ is the startup power profile and $[\cdot]_+$ caps its argument to positive values, i.e., $S_{u,i}(t) = 0$ for $t \geq N_{s,i}$.

If the power consumption of each refrigerator was constant over time, ΔD_t in (7) would be the fraction of loads that need to switch, and it would depend only on Δf_t and Δf_{t-1} . Since the startup dynamics affect the future aggregate power, the fraction of loads that need to switch at the current time step depends additionally on the number of loads that switched during the previous reserve activations. Therefore, the fraction of loads to be switched – denoted by x_t – is in general different from ΔD_t .

For control design we assume to know only the average values $\bar{P}_n = \mathbb{E}[P_{n,i}]$, $\bar{u} = \mathbb{E}[u_i]$ and $\bar{N}_s = \mathbb{E}[N_{s,i}]$ in the population⁴, and define the average startup duty cycle profile

$$\bar{S}_u(t) = \bar{u} \cdot \left[1 - \frac{t}{\bar{N}_s} \right]_+, \quad (9)$$

which we use to recursively calculate x_t with

$$x_t = \frac{1}{1 + \nu_t} \left[D_t^d - D_{t-1}^a - \sum_{i=t-N_s+1}^{t-1} x_i \bar{S}_u(t-i) \right] \quad (10)$$

$$D_t^a = D_{t-1}^a + x_t. \quad (11)$$

⁴The average values can be computed from available statistics on domestic refrigerators. Typical values are given in Table 7.

The variable D_t^a is the actual duty cycle of the population, whereas $\bar{S}_u(t)$ models the startup dynamics as the triangle shown in the left plot of Fig. 1. For $t = 0$, $\bar{S}_u(0)$ corresponds to the normalized overshoot in power consumption \bar{u} . For example, if the maximum power consumption of refrigerators at the startup phase is 25% higher than the nominal power, then $\bar{S}_u(0) = 0.25$. For $0 < t \leq N_s$, the variable $\bar{S}_u(t)$ models the decay of the originally higher power to the nominal power.

The term $1/(1+\nu_t)$ in (10) reduces the fraction of switched loads x_t as much as the expected overshoot due to the current startup dynamics. Note that $\nu_t = \bar{S}_u(0)$ if the right hand side of (10) is non-negative, whereas $\nu_t = 0$ if it is negative. This differentiation is necessary since there exist no dynamics at shutdown. On the other hand, the summation term of (10) recursively accounts for the anticipated power overshoot due to the past startup dynamics associated with recent switching actions up to time step $t - 1$.

Using (10) and (11) in the controller, we estimate the effect of refrigerators' startup dynamics on the aggregate power and the fraction of switched loads. Since (10) and (11) rely only on the mean values \bar{P}_n , \bar{u} and \bar{N}_s , the control will not be perfect. However, it is possible to characterize its performance analytically as shown by Proposition 1.

Proposition 1. *Let us denote by $N_{s,min}$ and $N_{s,max}$ the minimum and maximum startup duration in the population of refrigerators. The estimated aggregate power due to refrigerator startup dynamics obtained using the average values \bar{P}_n , \bar{u} , \bar{N}_s and (9) is an upper bound of the actual aggregate power up to $N_{s,min}$ seconds after the reserve activation, irrespective of the probability distribution of $N_{s,i}$.*

Furthermore, if $N_{s,i}$ follows the uniform distribution $N_{s,i} \sim [N_{s,min}, N_{s,max}]$, then the estimated aggregate power is an upper bound for up to t_{lim} seconds after the reserve activation, whereas it is a lower bound afterwards, where t_{lim} is given by

$$t_{lim} = \frac{N_{s,max} (N_{s,min} + N_{s,max})}{3N_{s,max} - N_{s,min}}. \quad (12)$$

Proof. The proof is given in the appendix. ■

Another effect of startup dynamics is an increase in the average power of the refrigerator aggregation, both with and without PFC. It is straightforward to express the increased aggregate power as a function of the average values of the parameters of startup dynamics (\bar{P}_n , \bar{u} , and \bar{N}_s), which is then used as a baseline for PFC evaluation purposes.

3.2 Compressor Lockout Constraints

Lockout times are usually employed to avoid compressor's frequent switching that decreases efficiency and possibly the lifetime [11]. Let us denote by t_{on}^l the lock-on time, i.e., the minimum duration the device must operate after switching on, and by t_{off}^l the lock-off time, i.e., the minimum duration the device must remain

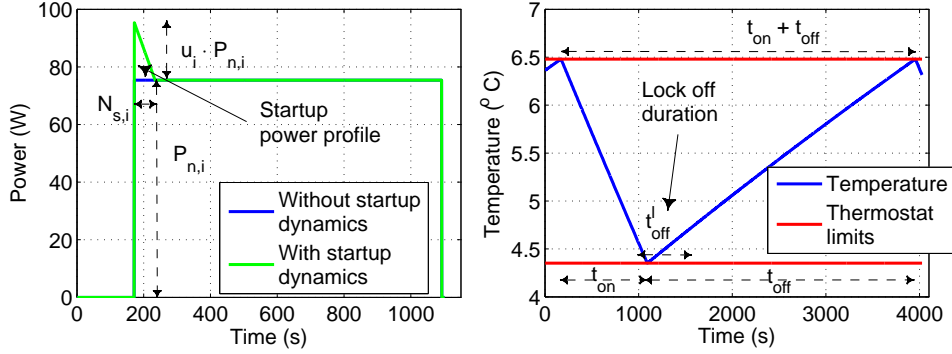


Figure 1: Left: Effect of startup dynamics on refrigerator power. Right: Evolution of refrigerator temperature during an on/off cycle.

off after switching off. When a compressor is locked it cannot react to frequency deviations, and so it is important to dynamically estimate the fraction of locked devices to avoid performance degradation or even loss of controllability. In the following, we present a decentralized method to do so.⁵

3.2.1 Steady-state Lockout due to Thermostatic Control Actions

Even without PFC, i.e., at steady state, at any time step t a fraction L_{on}^{st} of refrigerators is locked at the on state and a fraction L_{off}^{st} is locked at the off state due to thermostatic control actions. Considering the cycle of an individual refrigerator in Fig. 1, the lock event at the off state l_{off} can be modeled as the Bernoulli probability distribution

$$f_{off}^l(l_{off}) = \begin{cases} \frac{t_{off}^l}{t_{on} + t_{off}}, & \text{if } l_{off} = 1 \\ 1 - \frac{t_{off}^l}{t_{on} + t_{off}}, & \text{if } l_{off} = 0. \end{cases} \quad (13)$$

The fraction of refrigerators that is locked at the off state is defined and calculated using the properties of the Bernoulli PDF

$$L_{off}^{st} := \frac{1}{N_r} \cdot \sum_{i=1}^{N_r} l_{off,i} = \mathbb{E}[l_{off}] = \frac{t_{off}^l}{t_{on} + t_{off}}. \quad (14)$$

Since t_{off}^l , t_{on} and t_{off} are random variables, L_{off}^{st} is also a random variable with a so-called ratio distribution. Without knowledge of the PDF of t_{off}^l , t_{on} , t_{off} , the expectation of L_{off}^{st} can be approximated with a first order Taylor expansion [39]

$$\mathbb{E}[L_{off}^{st}] = \mathbb{E}\left[\frac{t_{off}^l}{t_{on} + t_{off}}\right] \approx \frac{\mathbb{E}[t_{off}^l]}{\mathbb{E}[t_{on}] + \mathbb{E}[t_{off}]}. \quad (15)$$

⁵The effect of lockout constraints on the refrigerators' aggregate response resembles a dead-band effect. However, a crucial difference is that the lockout effect is not memoryless, but it evolves in time depending on the statistics of lockout times. For this reason, the standard methods in control theory to handle the deadband effect cannot be used to model the effect of lockout constraints [38].

Using a similar procedure, one can get the approximation for steady-state lockout at the on state

$$\mathbb{E}[L_{\text{on}}^{\text{st}}] \approx \frac{\mathbb{E}[t_{\text{on}}^l]}{\mathbb{E}[t_{\text{on}}] + \mathbb{E}[t_{\text{off}}]} . \quad (16)$$

3.2.2 Transient Lockout due to Primary Frequency Control

When providing PFC the refrigerators are exposed to additional switching actions that increase the fraction of locked devices. Let us denote by $L_{\text{on},t}^{\text{tr}}$, $L_{\text{off},t}^{\text{tr}}$ the time-varying fractions of devices that lock at the on and off states due to PFC. Since $L_{\text{on},t}^{\text{tr}}$ and $L_{\text{off},t}^{\text{tr}}$ depend on the past reserve activations, they can be estimated recursively given sufficient statistical information. We assume that the Cumulative Distribution Functions (CDFs) of the lockout times are known, and we denote them by $F_{\text{on}}(t_{\text{on}}^l)$ for lock-on time and $F_{\text{off}}(t_{\text{off}}^l)$ for lock-off time. We further define $N_{\text{on}} = \max[t_{\text{on}}^l]$ and $N_{\text{off}} = \max[t_{\text{off}}^l]$. Then, $L_{\text{on},t}^{\text{tr}}$ and $L_{\text{off},t}^{\text{tr}}$ can be computed as

$$L_{\text{on},t}^{\text{tr}} = \sum_{k=0}^{t-1} c_k x_k S_{\text{on}}(t-k), \quad c_k = \begin{cases} 1, & \text{if } x_k \geq 0 \\ 0, & \text{if } x_k < 0 \end{cases} \quad (17)$$

$$L_{\text{off},t}^{\text{tr}} = \sum_{k=0}^{t-1} d_k x_k S_{\text{off}}(t-k), \quad d_k = \begin{cases} 0, & \text{if } x_k \geq 0 \\ 1, & \text{if } x_k < 0 \end{cases}, \quad (18)$$

where $S_{\text{on}} = 1 - F_{\text{on}}$ and $S_{\text{off}} = 1 - F_{\text{off}}$ are the survival functions of the CDFs, $S_{\text{on}}(t-k) = 0$ for $t-k > N_{\text{on}}$, and $S_{\text{off}}(t-k) = 0$ for $t-k > N_{\text{off}}$.

The total fraction of locked devices is simply obtained by adding the steady-state and transient contributions

$$L_{\text{on},t} = \mathbb{E}[L_{\text{on}}^{\text{st}}] + L_{\text{on},t}^{\text{tr}}, \quad L_{\text{off},t} = \mathbb{E}[L_{\text{off}}^{\text{st}}] + L_{\text{off},t}^{\text{tr}} . \quad (19)$$

3.3 Improved Probabilistic Switching

In Sections 3.1 and 3.2 we presented methods to account for the startup dynamics and lockout constraints when providing PFC with an aggregation of refrigerators. To integrate these methods in the probabilistic switching approach of Section 2.3, it suffices to apply the switching probability

$$\rho_t = \begin{cases} \frac{x_t}{1 - D_{t-1}^{\text{a}} - L_{\text{off},t-1}}, & \text{if } x_t \geq 0 \text{ (switch on)} \\ -\frac{x_t}{D_{t-1}^{\text{a}} - L_{\text{on},t-1}}, & \text{if } x_t < 0 \text{ (switch off)} \end{cases}, \quad (20)$$

instead of (7). Observe that the total fraction of loads that need to switch (x_t) is normalized by the fraction of loads that are available to respond. If $x_t > 0$, the loads that can respond are those that are at the off state and unlocked, i.e., the normalization factor is $1 - D_{t-1}^{\text{a}} - L_{\text{off},t-1}$. On the other hand, if $x_t < 0$, the available loads are those that are at the on state and unlocked, i.e., the normalization factor is $D_{t-1}^{\text{a}} - L_{\text{on},t-1}$.

In order to further reduce the compressor cycling, the switching probability can be defined as a function of the time elapsed since the latest switching action of each device. However, we do not use a time-dependent probability calculation in our controller, because it tends to synchronize the load cycles and introduce long-term oscillations in aggregate power.

4 Control of Thermostat Temperature Limits

4.1 Resetting of Thermostat Limits

The probabilistic switching of Section 3.3 provides accurate PFC reserves immediately after a frequency deviation. However, if the frequency deviation remains, the accuracy of reserve provision will deteriorate because the refrigerators' aggregate power will decay towards the pre-disturbance value. This is straightforward to verify if the startup dynamics and lockout constraints are neglected.

Let $s_{\text{off},t}$ and $s_{\text{on},t}$ denote the rates at which refrigerators switch off and on at time step t , which we approximate with

$$s_{\text{off},t} \approx D_t^d \cdot \dot{T}_d(\bar{T}), \quad s_{\text{on},t} \approx (1 - D_t^d) \cdot \dot{T}_i(\bar{T}), \quad (21)$$

where $\dot{T}_d(\bar{T})$ and $\dot{T}_i(\bar{T})$ are the temperature decrease and increase rates evaluated at the mean temperature, respectively. $\dot{T}_d(\bar{T})$ can be obtained from (2) and $\dot{T}_i(\bar{T})$ from (1), and for the average values of the parameters used in this paper we get $\dot{T}_d(\bar{T}) = -0.0026^\circ\text{C/s}$ and $\dot{T}_i(\bar{T}) = 0.0009^\circ\text{C/s}$. Combining (21) and (6), the net switching rate can be expressed as

$$\begin{aligned} s_{\text{net},t} &= s_{\text{off},t} + s_{\text{on},t} \\ &\approx D^n \dot{T}_d(\bar{T}) + (1 - D^n) \dot{T}_i(\bar{T}) + D^r \frac{\Delta f_t}{\Delta f_{\text{max}}} [\dot{T}_d(\bar{T}) - \dot{T}_i(\bar{T})] \\ &= D^r \cdot (\Delta f_t / \Delta f_{\text{max}}) \cdot [\dot{T}_d(\bar{T}) - \dot{T}_i(\bar{T})]. \end{aligned} \quad (22)$$

In (22) $s_{\text{off}} = D^n \cdot \dot{T}_d(\bar{T})$ and $s_{\text{on}} = (1 - D^n) \cdot \dot{T}_i(\bar{T})$ are the switching rates at steady state and $s_{\text{off}} + s_{\text{on}} = 0$ holds because the aggregate baseline power is approximately constant.

If $\Delta f_t > 0$, some refrigerators will switch on to provide PFC and $s_{\text{net},t}$ will become negative because $\dot{T}_d(\bar{T}) - \dot{T}_i(\bar{T}) < 0$. This means that some refrigerators will start switching off and therefore the aggregate power will start decreasing towards the baseline. Similarly, if $\Delta f_t < 0$ the initially lower aggregate power will start increasing towards the baseline. Thus, the reserve accuracy will deteriorate over time.

A solution to this problem is to dynamically modify the devices' thermostat limits in order to keep $s_{\text{net},t}$ close to zero. For this purpose, we apply the temperature resetting factor

$$K_{r,t} [^\circ\text{C}] = s_{\text{net},t} \cdot \Delta t = -D^r \Delta t \bar{\beta} \bar{P}_n \cdot \frac{\Delta f_t}{\Delta f_{\text{max}}}, \quad (23)$$

where $\dot{T}_d(\bar{T}) - \dot{T}_i(\bar{T}) = -\bar{\beta}\bar{P}_n$ from (1) and (2). If $\Delta f_t > 0$ the resetting factor is negative, which means that both thermostat limits will decrease with a rate equal to $K_{r,t}$ such that the population reaches an equilibrium at an aggregate power higher than the baseline. Similarly, if $\Delta f_t < 0$ both thermostat limits increase at a rate equal to $K_{r,t}$.

To account for startup dynamics and lockout constraints a few modifications are needed in (23). The first modification is to substitute the term $D^r \cdot (\Delta f_t / \Delta f_{\max})$ with x_t to account for startup dynamics – see (11). The second modification is to incorporate the sequential unlocking of refrigerators after a frequency deviation event based on the CDFs of lockout times $F_{\text{on}}(t_{\text{on}}^l)$ and $F_{\text{off}}(t_{\text{off}}^l)$.

Consider a positive Δf_t starting at time step $t = 0$ that will induce a positive x_t . Instantaneously, some refrigerators will switch on to provide PFC and will lock at the on state. Therefore, at $t = 0$ fewer devices will be at the off state, $s_{\text{on},t}$ will start decreasing, and an instantaneous resetting factor of $-x_t \dot{T}_i(\bar{T})$ is needed to keep $s_{\text{net},t}$ close to zero. However, as time elapses, some of the refrigerators that switched on will unlock based on $F_{\text{on}}(t_{\text{on}}^l)$, and larger changes in the temperature deadband limits are required to maintain the same aggregate power. Therefore, for $\Delta f_t > 0$ the resetting factor is determined by $x_t [\dot{T}_d(\bar{T}) F_{\text{on}}(t) - \dot{T}_i(\bar{T})] \Delta t$ and for $\Delta f_t < 0$ by $x_t [\dot{T}_d(\bar{T}) - \dot{T}_i(\bar{T}) F_{\text{off}}(t)] \Delta t$.

To generalize for multiple events, the resetting factor at time step t should accumulate all contributions from the previous time steps. Therefore, we compute the resetting factor with

$$K_{r,t} = r_t \cdot \Delta t \cdot \sum_{k=0}^{t-1} x_k \cdot [c_k \cdot [\dot{T}_d(\bar{T}) \cdot F_{\text{on}}(t-k) - \dot{T}_i(\bar{T})] + d_k \cdot [\dot{T}_d(\bar{T}) - \dot{T}_i(\bar{T}) \cdot F_{\text{off}}(t-k)]], \quad (24)$$

where c_k and d_k are defined in (17) and (18). Recall that $F_{\text{on}}, F_{\text{off}}$ are CDFs, and so $F_{\text{on}}(t-k) = 1$ for $t-k > N_{\text{on}}$ and $F_{\text{off}}(t-k) = 1$ for $t-k > N_{\text{off}}$ in (24). Note that the temperature resetting is applied only to the unlocked refrigerators, because only those affect the net switching rate. This is accounted for in (24) with the normalization factor

$$r_t = \frac{1 - \mathbb{E}[L_{\text{on}}^{\text{st}}] - \mathbb{E}[L_{\text{off}}^{\text{st}}]}{1 - L_{\text{on},t} - L_{\text{off},t}}. \quad (25)$$

Observe that $r_t > 1$ as long as $L_{\text{on},t}^{\text{tr}} \neq 0$ or $L_{\text{off},t}^{\text{tr}} \neq 0$, whereas $r_t = 1$ if $L_{\text{on},t}^{\text{tr}} = L_{\text{off},t}^{\text{tr}} = 0$, i.e., at steady state.

By modifying the thermostat limits using (24) we effectively transform the electric energy surplus or deficit due to the PFC reserve into thermal energy. In this way, we avoid the decay of aggregate power to the pre-disturbance value and achieve constant reserve provision.

4.2 Corrective Temperature Control

If the frequency deviation Δf_t is approximately zero-mean, the population's mean temperature does not change significantly and the duty cycle D_t remains close to

the nominal value D^n . However, in case of biased frequency deviations the resetting factor $K_{r,t}$ results in prolonged reduction or increase of the thermostat limits, which is undesirable for two main reasons. First, it negatively affects the user utility due to very high or very low temperatures. Second, the population's mean temperature changes substantially, which results in steady-state baseline deviations that affect the generation-demand balance on the system level. A simple solution to this problem would be to filter the frequency deviation signal in order to eliminate the bias [31]. However, the disadvantage of this approach is that the bias needs to be absorbed by other available PFC resources or transferred to secondary frequency control.

We propose a different approach that allows us to handle biased frequency deviations without relying on additional resources. An opposite control action is imposed on the thermostat limits such that they return to the nominal values in the long run. This “*corrective temperature control*” is designed as a proportional feedback controller on the average temperature across the population using a gain K_c . Thus, the temperature limits of refrigerator i evolve in time according to

$$T_{\min,i,t} = T_{\min,i,t-1} + K_{r,t} - K_c \cdot (\bar{T}_{t-1} - \bar{T}_{\text{nom}}) \quad (26)$$

$$T_{\max,i,t} = T_{\max,i,t-1} + K_{r,t} - K_c \cdot (\bar{T}_{t-1} - \bar{T}_{\text{nom}}), \quad (27)$$

where \bar{T}_{t-1} is an estimate of the population's average temperature, and \bar{T}_{nom} is the nominal mean temperature without PFC.

4.2.1 Estimation of Mean Temperature and Duty Cycle

The mean temperature \bar{T}_t depends on the past temperature resetting and corrective control actions according to

$$\bar{T}_t = \bar{T}_{t-1} + [K_{r,t} - K_c(\bar{T}_{t-1} - \bar{T}_{\text{nom}})] \cdot (1 - L_{\text{on},t} - L_{\text{off},t}). \quad (28)$$

The term $(1 - L_{\text{on},t} - L_{\text{off},t})$ in (28) is needed because the temperature resetting and corrective control are applied only to the unlocked devices. With reference to (15) and (16), the fractions of loads that are locked at steady-state depend on the duration of the on and off cycles. However, as the mean temperature changes, the expected values of the on and off cycles also change and are approximated with

$$\mathbb{E}[t_{\text{on},t}] \approx \bar{R}\bar{C} \cdot \ln \left[\frac{\bar{T}_{\max,t-1} - \bar{T}_a + \bar{\eta}\bar{R}\bar{P}_n}{\bar{T}_{\min,t-1} - \bar{T}_a + \bar{\eta}\bar{R}\bar{P}_n} \right] \quad (29)$$

$$\mathbb{E}[t_{\text{off},t}] \approx \bar{R}\bar{C} \cdot \ln \left[\frac{\bar{T}_a - \bar{T}_{\min,t-1}}{\bar{T}_a - \bar{T}_{\max,t-1}} \right], \quad (30)$$

where $\bar{T}_{\min,t-1} = \bar{T}_{t-1} - 0.5 \cdot \bar{\Delta T}$, $\bar{T}_{\max,t-1} = \bar{T}_{t-1} + 0.5 \cdot \bar{\Delta T}$, and $\bar{\Delta T}$ is the deadband width.

Using \bar{T}_t it is also possible to keep track of the population's baseline duty cycle (without PFC) with the approximation

$$\bar{D}_t^n \approx \frac{\mathbb{E}[t_{\text{on},t}]}{\mathbb{E}[t_{\text{on},t}] + \mathbb{E}[t_{\text{off},t}]} . \quad (31)$$

The actual duty cycle while providing PFC is given by

$$D_t^a = D_{t-1}^a + x_t + (\bar{D}_t^n - \bar{D}_{t-1}^n) , \quad (32)$$

where x_t is the fraction of loads activated for PFC, and $\bar{D}_t^n - \bar{D}_{t-1}^n$ is the change in the baseline duty cycle due to deviation of the mean temperature from its nominal value.

Although the average temperature can be estimated in a decentralized way using (28), the estimates could be reset to the actual values periodically based on measurements of refrigerator temperatures transmitted via a low data transfer communication link (see Section 2.4). Apart from providing an initial temperature estimate when the aggregation starts providing PFC reserves, a communication link will also enhance the controller's robustness to estimation error.

4.2.2 Analytical Tuning of the Corrective Temperature Gain

From a user point of view, high K_c values are preferable to keep the temperature deviations low. However, from a power system point of view, there is a tradeoff to consider when tuning K_c . If the gain is very small, prolonged temperature deviations will change the aggregation's baseline and introduce steady-state reserve errors. On the other hand, a very large gain might introduce oscillations in the aggregate power and thus increased reserve errors. In fact, upper and lower bounds on K_c can be computed analytically.

The correction of mean temperature with a gain K_c can be thought as temperature resetting with the $K_{r,t}$ defined in (24), but for an opposite frequency deviation signal. Assuming that the application of the temperature resetting factor $K_{r,t}$ does not create oscillations in aggregate power, and if the opposite temperature change due to K_c is in absolute terms smaller than the average $\bar{K}_{r,t}$, then no oscillations in aggregate power should occur. Therefore, an upper bound on K_c can be obtained from

$$\begin{aligned} |K_c \cdot (\bar{T}_t - \bar{T}_{\text{nom}})| &\leq |\bar{K}_{r,t}| \Leftrightarrow \\ K_c \cdot |\bar{T}_t - \bar{T}_{\text{nom}}| &\leq \Delta t \bar{\beta} \bar{P}_n \left| \frac{D^r \Delta f_t}{\Delta f_{\text{max}}} \right| = \Delta t \bar{\beta} \bar{P}_n \cdot |\bar{D}_t^n - D^n|. \end{aligned} \quad (33)$$

From (31) we can see that \bar{D}_t^n is a function of the mean temperature \bar{T}_t . Therefore, the upper bound on K_c from (33) is in general a function of \bar{T}_t itself. However, a temperature-independent bound can be obtained by a first order Taylor expansion of (31) that gives

$$\bar{D}_t^n(\bar{T}) \approx \bar{D}_t^n(\bar{T}_{\text{nom}}) + \left. \frac{d\bar{D}_t^n(\bar{T})}{d\bar{T}} \right|_{\bar{T}_{\text{nom}}} \cdot (\bar{T} - \bar{T}_{\text{nom}}) \quad \Leftrightarrow$$

$$\bar{D}_t^n - D^n = \bar{D}_t^n(\bar{T}) - \bar{D}_t^n(\bar{T}_{\text{nom}}) \approx \left. \frac{d\bar{D}_t^n(\bar{T})}{d\bar{T}} \right|_{\bar{T}_{\text{nom}}} \cdot (\bar{T} - \bar{T}_{\text{nom}}). \quad (34)$$

Substituting now (34) in (33) gives us the upper bound

$$K_c \leq \left| \Delta t \bar{\beta} \bar{P}_n \cdot \left. \frac{d\bar{D}_t^n(\bar{T})}{d\bar{T}} \right|_{\bar{T}_{\text{nom}}} \right|. \quad (35)$$

A lower bound on K_c can be obtained based on the minimum/maximum acceptable mean temperature and the maximum acceptable settling time, i.e., the time needed to restore \bar{T}_t close enough to \bar{T}_{nom} after the biased frequency deviation event is over. Neglecting lockout constraints (28) can be written as

$$\bar{T}_t = \lambda \cdot \bar{T}_{t-1} - \gamma \cdot \Delta f_t + (1 - \lambda) \cdot \bar{T}_{\text{nom}}, \quad (36)$$

where $\lambda = 1 - K_c$ and $\gamma = D^r \Delta t \bar{\beta} \bar{P}_n / \Delta f_{\text{max}}$. With an initial condition \bar{T}_0 , (36) has the solution

$$\bar{T}_t = \lambda^t \bar{T}_0 - \gamma \sum_{k=0}^{t-1} \lambda^k \Delta f_{t-k-1} + (1 - \lambda^t) \bar{T}_{\text{nom}}. \quad (37)$$

Assume that $\bar{T}_0 = \bar{T}_{\text{nom}}$ and that the refrigerator aggregation faces the step frequency deviation

$$\Delta f_t = \begin{cases} \delta, & \text{if } 0 \leq t \leq N_{\text{ev}} \\ 0, & \text{if } t > N_{\text{ev}} \end{cases}, \quad (38)$$

which represents a frequency deviation signal with a bias equal to δ over a period N_{ev} . Let ε denote the tolerance in terms of mean temperature during the frequency deviation event, i.e., \bar{T}_t must satisfy $\bar{T}_{\text{nom}} - \varepsilon \leq \bar{T}_t \leq \bar{T}_{\text{nom}} + \varepsilon \forall t$. Furthermore, let N_{rec} denote the maximum acceptable settling time after the frequency deviation event, i.e., $\bar{T}_{\text{nom}} - \epsilon \leq \bar{T}_t \leq \bar{T}_{\text{nom}} + \epsilon$ should hold for $t \geq N_{\text{ev}} + N_{\text{rec}}$, where $\epsilon < \varepsilon$ is another tolerance. A lower bound on K_c can be computed using Proposition 2.

Proposition 2. *If $\bar{T}_0 = \bar{T}_{\text{nom}}$ and Δf_t is given by (38), then $\bar{T}_{\text{nom}} - \varepsilon \leq \bar{T}_t \leq \bar{T}_{\text{nom}} + \varepsilon \forall t$ and $\bar{T}_{\text{nom}} - \epsilon \leq \bar{T}_t \leq \bar{T}_{\text{nom}} + \epsilon$, for $t \geq N_{\text{ev}} + N_{\text{rec}}$, hold if the corrective temperature gain is computed as $K_c = 1 - \lambda$, where λ satisfies*

$$\gamma \delta \cdot (1 - \lambda^{N_{\text{ev}}}) \leq \varepsilon \cdot (1 - \lambda) \quad (39)$$

$$\gamma \delta \cdot \lambda^{N_{\text{rec}}} \cdot (1 - \lambda^{N_{\text{ev}}}) \leq \epsilon \cdot (1 - \lambda). \quad (40)$$

Proof. The proof is given in the appendix. ■

Remark 1. *Although (39), (40) cannot be solved analytically, a lower bound on K_c is obtained numerically starting with $\lambda = 1$ and gradually reducing it until (39), (40) are satisfied.*

Equation (35), Proposition 2, and Remark 1 provide theoretical lower and upper bounds on K_c that can serve as an initial range when determining the final value of the gain via simulations with realistic Δf_t signals, and in the presence of startup dynamics and lockout constraints.

4.3 Addressing Limited Thermostat Resolution

Typically, the change in the thermostat's temperature limits from (26) and (27) is very small, for example, 10^{-4} °C/s. Nevertheless, the measurement accuracy in a refrigerator can be as low as ± 0.2 °C [40]. Therefore, the temperature limits of some refrigerators might not change, and thus the aggregate power might decay towards its pre-disturbance value.

We address this practical issue by allowing the temperature limits to change according to a sufficiently large fixed rate ΔT_{res} °C per second. Specifically, instead of distributing the temperature change over the whole population and requesting a small change from every device, we impose a larger change on the temperature limits of fewer devices.

4.3.1 Probabilistic Implementation

The required change in temperature limits is

$$\Delta T_{\text{lim},t} = K_{r,t} - K_c \cdot (\bar{T}_{t-1} - \bar{T}_{\text{nom}}) . \quad (41)$$

If the temperature limits of an individual device change only by ΔT_{res} , we can get the same total change in temperature by applying ΔT_{res} to a fraction of the population equal to

$$\varrho_t = |\Delta T_{\text{lim},t}| / \Delta T_{\text{res}} . \quad (42)$$

The implementation of this approach is straightforward: each refrigerator draws a random number uniformly distributed between 0 and 1 and if it is smaller than ϱ_t and the refrigerator is unlocked, then it modifies its temperature limits by ΔT_{res} , otherwise it does not. The limits increase by ΔT_{res} if $\Delta T_{\text{lim},t} > 0$ and decrease by ΔT_{res} if $\Delta T_{\text{lim},t} < 0$.

Although simple, this approach introduces two sources of inaccuracy in the control. First, as ΔT_{res} increases, fewer devices change their thermostat limits and therefore the control is exposed to inaccuracies from random number generation. Second, large ΔT_{res} values might synchronize the population. Without the fixed rate ΔT_{res} , the temperature limits change in a continuous way, the devices switch when they are close to the thermostat limits and this avoids synchronization. In contrast, if a large ΔT_{res} value is used, some devices might switch at a temperature considerably different than the thermostat limits.

4.3.2 Effect on Temperature Limits and Countermeasure

If a fixed rate ΔT_{res} is used, the change in temperature limits $\Delta T_{\text{lim},i,t}$ is different for each refrigerator i . A disadvantage of this approach is that the variance of $\Delta T_{\text{lim},i,t}$ monotonically increases, as shown by Proposition 3 and Remark 2.

Proposition 3. *If the refrigerators' temperature limits change probabilistically according to (42), then the mean and variance of $\Delta T_{lim,i,t}$ are given by*

$$\mathbb{E} [\Delta T_{lim,i,t}] = \sum_{k=0}^{t-1} \Delta T_{lim,k} \quad (43)$$

$$\text{Var} [\Delta T_{lim,i,t}] = \sum_{k=0}^{t-1} |\Delta T_{lim,k}| \cdot (\Delta T_{res} - |\Delta T_{lim,k}|). \quad (44)$$

Proof. The proof is given in the appendix. ■

Remark 2. *If no fixed rate ΔT_{res} is applied, all refrigerators change their temperature limits by $\sum_{k=0}^{t-1} \Delta T_{lim,k}$, so the mean value is $\sum_{k=0}^{t-1} \Delta T_{lim,k}$ and the variance is zero. In practice, the mean value will be slightly different and the variance non-zero but small because some devices are always locked at steady-state and do not change their temperature limits. According to (43), ΔT_{res} does not affect the mean value of temperature limit change. In contrast, the variance monotonically increases in time (linearly with ΔT_{res}) because $\Delta T_{res} - |\Delta T_{lim,k}| \geq 0$.*

Due to the monotonic increase in variance, some devices will likely sustain larger and prolonged deviations from the nominal temperature limits, which is undesirable. The variance can be bounded by imposing a bound ΔT_b on the maximum deviation from the mean temperature of the population. To implement this, we require each unlocked device i that drew a random number smaller than ϱ_t at time step t to change its limits only if

$$\bar{T}_{t-1} - \Delta T_b \leq \Delta T_{lim,i,t-1} + \zeta_i \Delta T_{res} \leq \bar{T}_{t-1} + \Delta T_b, \quad (45)$$

where ζ_i is equal to 1 or -1 depending on the sign of $\Delta T_{lim,t}$.

5 Refrigerator Door Openings

The analysis so far neglected the door openings that significantly affect the aggregate power of a population of refrigerators. Since the PFC performance is evaluated with respect to a baseline, neglecting the effect of door openings will introduce large control errors. In this section, we present a method to account for door openings in PFC design.

5.1 Stochastic Model for Door Opening Events

There exist only a few papers that consider door openings in power system studies with refrigerators, e.g., [10, 27]. We assume that the number of door openings per day follows a normal distribution with mean value μ_{op} and standard deviation σ_{op} . Furthermore, we assume that the duration of a door opening event follows a normal distribution with mean value μ_d and standard deviation σ_d . Based on [41] we select $\mu_{op} = 40$, whereas we fix $\sigma_{op} = 5$, $\mu_d = 20$ s and $\sigma_d = 3$ s [27].

5.2 Modeling the Effect of Door Openings

The door openings increase the refrigerator's energy consumption by $\xi \cdot 100\%$ (we assume $\xi = 0.22$ [41]). The effect of door openings on refrigerator temperature is modeled by reducing the thermal resistance of the model (1), (2) during each door opening event, i.e., by increasing the thermal losses to the ambient. An upper bound to the new thermal resistance is given by Proposition 4.

Proposition 4. *Denote by $N^d = 86,400$ s the duration of a day. An upper bound to the estimate of the thermal resistance R_{op} during a door opening event can be obtained from*

$$R_{op} \leq R \cdot \frac{1}{1 + \frac{N^d}{\mu_{op} \cdot \mu_d} \cdot \xi} . \quad (46)$$

Proof. The proof is given in the appendix. ■

Remark 3. *Evaluating (46) as equality gives us an initial estimate of R_{op} . For the assumed parameters, (46) gives $R_{op} \leq R/24.76$, and so we set $R_{op} = R/25$ in our simulations.*

5.3 Modifications in Primary Frequency Control Design

The proposed controller can be easily extended to account for door openings. Even without PFC, the aggregation's duty cycle and baseline power are not constant any more, but they depend on the distribution of door openings within the day. The duty cycle without PFC can be expressed as

$$D_t^n = D^n + \Delta D_t^{op} , \quad (47)$$

where $\Delta D_t^{op} \geq 0$ is the additional duty cycle due to door openings. A smoothed version of ΔD_t^{op} can be obtained comparing historical data of aggregate power consumption with and without door openings.

Instead of (31), the duty cycle is now estimated using

$$\bar{D}_t^n \approx \frac{\mathbb{E}[t_{on,t}]}{\mathbb{E}[t_{on,t}] + \mathbb{E}[t_{off,t}]} + \Delta D_t^{op} . \quad (48)$$

The necessary information about door openings is passed to the switching probability calculation (20) through (32), and no other modification in the control design is needed. In particular, note that each refrigerator responds to frequency deviations irrespective of whether the door is closed or open.

6 Parameters and Performance Metrics

The proposed controller (consisting of all components presented in the previous sections) is graphically shown in Fig. 2 and is easy to implement at the device

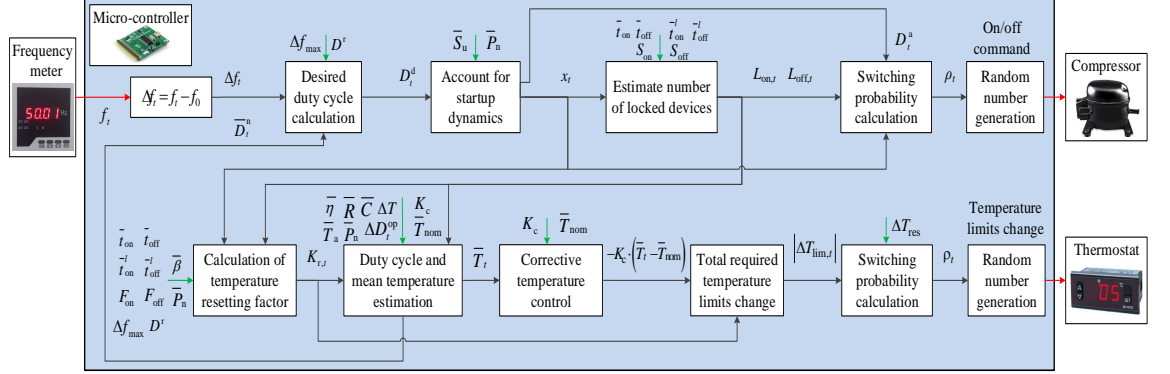


Figure 2: The proposed decentralized stochastic controller. Red arrows indicate the controller's input/output signals, green arrows indicate input from stored parameter values in the controller's memory, and black arrows indicate internal control and feedback signals.

Table 7: Refrigeration Aggregation Parameters

Parameter	Value	Parameter	Value
T_a (°C)	$\mathbb{U}[20, 24]$	α (s ⁻¹)	$\mathbb{U}[4, 6] \cdot 10^{-5}$
ΔT (°C)	$\mathbb{U}[1.7, 2.3]$	β (°C/kJ)	$(\mu, \sigma) = (4.4, 0.7) \cdot 10^{-5}$
T_{set} (°C)	$\mathbb{U}[4.5, 5.5]$	P_n (W)	$\mathbb{U}[70, 90]$
u (-)	$\mathbb{N}[0.25, 0.025]$	t_{on}^l (s)	$\mathbb{N}[60, 5]$
N_s (s)	$\mathbb{N}[30, 3]$	t_{off}^l (s)	$(\mu, \sigma) = (189, 31.5)$

level. The required hardware includes a frequency meter and a micro-controller to control the switch of the compressor and the temperature limits of the thermostat. The micro-controller consists of a random number generator and a floating-point unit able to perform elementary mathematical operations (additions, subtractions, multiplications and divisions).

In this section, we present the refrigerator parameters in Table 7, and introduce the metrics and benchmarks that we use to quantify the controller's performance. Note that β and t_{off}^l have the mean value and standard deviation shown in Table 7, but they do not follow a standard PDF because they are post-calculated based on the imposed PDFs on other variables [36].

For a fixed D^f the PFC reserve capacity is given by

$$P_{\text{res}} = N_r \bar{P}_n D^f. \quad (49)$$

The desired aggregate power at time step t with PFC is

$$P_{d,t} = \bar{P}_{b,t} + P_{\text{res}} \cdot (\Delta f_t / \Delta f_{\text{max}}), \quad (50)$$

where $\bar{P}_{b,t}$ is the population's baseline power, i.e., a smoothed version of historical measurements of aggregate power $P_{b,t}$ without PFC. If $w_i(t) = 0 \forall i$, the baseline is constant and equal to $\bar{P}_b = N_r \bar{P}_n D^n$.

The reserve Mean Absolute Percentage Error (MAPE) $e_{r,\text{mape}}$ and the tracking MAPE $e_{t,\text{mape}}$ are used as control performance metrics. Note that in the definition of tracking MAPE the normalized error is computed by dividing the absolute error with the desired power, whereas in the reserve MAPE definition the normalization is performed using the reserve capacity. In addition, we use the baseline MAPE $e_{b,\text{mape}}$ as a metric of the natural oscillations in an uncontrolled refrigerator aggregation. The metrics are defined as (N_{sim} is the simulation period)

$$e_{r,\text{mape}} = (100/N_{\text{sim}}) \cdot \sum_{t=0}^{N_{\text{sim}}-1} |(P_{d,t} - P_{\text{agg},t})/P_{\text{res}}| \quad (51)$$

$$e_{t,\text{mape}} = (100/N_{\text{sim}}) \cdot \sum_{t=0}^{N_{\text{sim}}-1} |(P_{d,t} - P_{\text{agg},t})/P_{d,t}| \quad (52)$$

$$e_{b,\text{mape}} = (100/N_{\text{sim}}) \cdot \sum_{t=0}^{N_{\text{sim}}-1} |(\bar{P}_{b,t} - P_{b,t})/P_{\text{res}}|. \quad (53)$$

The “*proposed controller*” is benchmarked against a “*simple controller 1*” that neglects the startup dynamics and lockout constraints, and is similar to the approach of [31]. The only difference is that the temperature resetting factor $K_{r,t}$ is computed analytically from (24) and not chosen heuristically as in [31]. We also consider a “*simple controller 2*” that apart from neglecting the startup dynamics and lockouts, it does not reset the thermostat limits as well.

The controllers' comparison is performed via detailed simulations using actual frequency deviation data from Switzerland in 2009 and in 2011. We assume that the refrigerators participate with a small share in PFC, and therefore have a negligible effect on system frequency. For this reason, it is sufficient to model the frequency as a time series in our simulations in order to evaluate the controller's performance. Of course, a dynamic frequency model incorporating also the generator's power response is needed for high shares of refrigerators in PFC, because their effect on frequency cannot be neglected anymore. In [42] we present results from dynamic frequency studies to show the contribution of refrigerators in reducing under-frequency after sudden loss of generation.

7 Sensitivity to Design and Load Parameters

7.1 Sensitivity to Aggregation Size

The control is expected to perform better for large refrigerator aggregations for three main reasons. First, the larger the aggregation the closer the fraction of switched loads to the desired switching probability. Second, the controller is designed based on average refrigerator parameters. And third, the assumption of a




Figure 3: Dependence of control performance and baseline error on the size of refrigerator aggregation for the proposed controller and simple controller 1.

smooth uncontrolled duty cycle D_t^n is reasonable only for sufficiently large aggregations, where the oscillations from individual load cycles cancel out.

We investigate the dependence of controller performance on aggregation size by running simulations using 30 zero-mean, 5-hour samples of frequency deviation ($D^r = 0.15$ is used). Figure 3 shows the reserve MAPE ($e_{r,mape}$) for the proposed and the simple controller 1, which decrease with the square root of the aggregation size. A population of 1,000 refrigerators results in a large reserve error, which drastically decreases when the size increases to 10,000. After this point, the reserve error decreases asymptotically and it practically saturates at an aggregation size of 70,000. The proposed controller outperforms the simple controller 1 for all aggregation sizes apart from an aggregation of 2,000 refrigerators. For populations of 70,000 loads or more, the proposed controller reduces the reserve error by 15%.

Simple controller 1 might perform better than the proposed controller for small aggregations, due to the dominant effect of the baseline MAPE ($e_{b,mape}$) for aggregations up to 15,000 loads, as shown in Fig. 3. In other words, the PFC response of small aggregations is covered by the large natural oscillations of the uncontrolled baseline. Therefore, the improvement of the proposed controller over simple controller 1 is guaranteed for sufficiently large aggregations, but not for small ones. Interestingly, the baseline error comprises a significant part of the control error even for large aggregations (for example, approximately 62% of the total error for 100,000 loads).

7.2 Sensitivity to Reserve Capacity

We investigate the controller's sensitivity to reserve capacity D^r for an aggregation of 70,000 refrigerators (according to the results of Fig. 3). We ran simulations for




Figure 4: Sensitivity to reserve capacity in simulations with small frequency deviation bias. Left: Reserve MAPE. Right: Tracking MAPE.

10 days with different values of D^r using as base case scenario the parameters of Table 7. The frequency deviation is approximately zero-mean during the first 5 days of the data set, whereas it has a significant bias during the remaining 5 days. The same frequency data are used for the all analyses of the rest of Section 7.

Figure 4 shows the sensitivity analysis results for the frequency deviations with small bias, whereas Fig. 5 is for the frequency deviations with large bias ($K_c = 0.5 \cdot 10^{-4}$ is used based on the controller tuning results of Section 8.1). In both figures, the left plot shows the reserve MAPE and the right plot the tracking MAPE.

The reserve capacity that minimizes the control error depends on the definition of the error metric. Specifically, as D^r increases the tracking MAPE increases but the reserve MAPE decreases. Therefore, although increasing the reserve capacity generally results in a higher control error, the error becomes smaller in relative terms (i.e., in comparison with the reserve capacity that is offered). The proposed controller consistently reduces the errors compared with simple controller 1. In addition, the latter performs better than simple controller 2, which illustrates the need for resetting of thermostat temperature limits.

Observe that the difference between the proposed controller and simple controller 1 is more pronounced for biased frequency deviations. This shows the effectiveness of the corrective temperature control of Section 4.2. In fact, there is only a small difference between the simple controller 1 and simple controller 2 for biased frequency deviations, because in both cases most of the reserve error is due to steady-state baseline errors.

7.3 Sensitivity to Peak Power of Startup Dynamics

The dependence of reserve MAPE on the peak power factor (u) during the startup phase is presented in Fig. 6. The error of the proposed controller is not sensitive to u , whereas the error of the simple controller 1 clearly increases as u increases. This is expected because larger u values result in larger overshoots in aggregate




Figure 5: Sensitivity to reserve capacity in simulations with large frequency deviation bias. Left: Reserve MAPE. Right: Tracking MAPE.




Figure 6: Sensitivity of reserve MAPE to peak power during startup dynamics. Left: Small frequency deviation bias. Right: Large frequency deviation bias.

power consumption, if the effect of startup dynamics is not considered in the control design.

7.4 Sensitivity to Lockout Times

Figure 7 shows the dependence of reserve error to the lock-on time t_{on}^l . The error of the proposed controller steadily increases as t_{on}^l increases, which indicates that the estimates of the number of locked devices – obtained by (19) – deteriorate. The error of simple controller 1 is significantly higher than that of the proposed controller, but is practically unaffected by t_{on}^l . This happens because simple controller 1 neglects both the lockouts and startup dynamics. With reference to (20), simple controller 1 activates fewer loads for PFC than actually needed, because it neglects the lockouts. However, each of the activated loads consumes more power than assumed by the controller due to startup dynamics, which counteracts the error in the




Figure 7: Sensitivity of reserve MAPE to lock-on time. Left: Small frequency deviation bias. Right: Large frequency deviation bias.

number of activated loads.

7.5 Effect of PFC Activation Deadband

In practice, a frequency deadband is used to prevent PFC activation to the wrong direction due to frequency measurement error. To investigate the effect of frequency deadband on the performance of the proposed controller, we repeated the simulations for our base case scenario with the deadband ± 10 mHz of the ENTSO-E system [43].

Even with a deadband, the proposed controller significantly reduces the reserve error compared with simple controller 1; however, the difference between the controllers is smaller because the PFC is active only for approximately 60% of the time. Figure 8 shows the desired and activated reserve for the two controllers with a PFC deadband for small and large frequency deviation bias. The proposed controller follows the desired droop characteristic of PFC much closer than simple controller 1, in particular for biased frequency deviations where the error is reduced from 4.3% to 1.2%.

7.6 Evaluation

Table 8 summarizes the average values of reserve MAPE computed across all simulation sets. The reduction in control error achieved by the proposed controller (in comparison with the simple controllers) depends strongly on the frequency deviation bias. If the frequency deviation is zero-mean, the error reduces by approximately 15%. For small frequency deviation bias, the proposed controller reduces the error by approximately 56% compared with simple controller 1 and 74% compared with simple controller 2. For large frequency deviation bias, the error reduction is approximately 78% compared with simple controller 1 and 82% compared with simple controller 2.

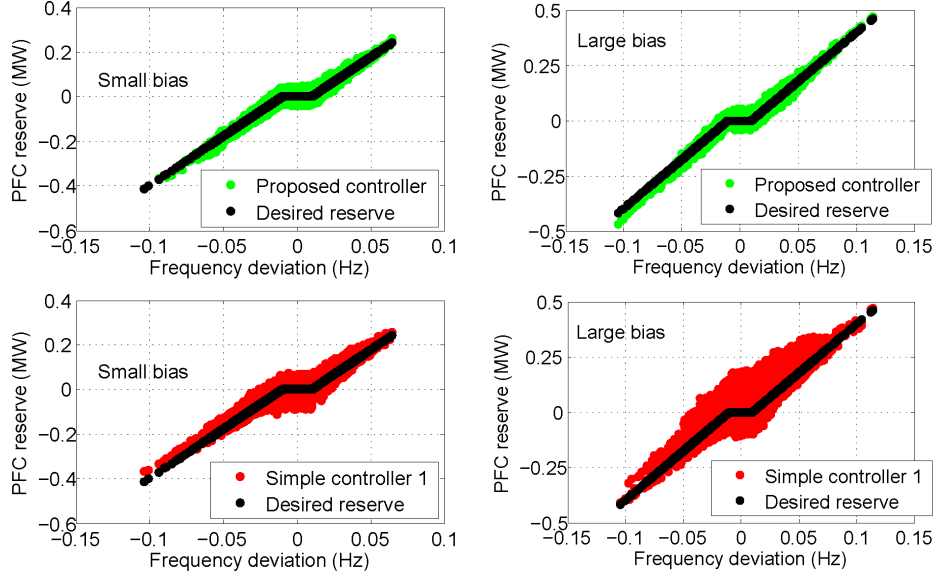


Figure 8: Droop characteristic with a PFC activation deadband. Left: Small frequency deviation bias. Right: Large frequency deviation bias.

Table 8: Average Reserve MAPE (%) for Different Frequency Biases

Controller	Zero-mean	Small bias	Large bias
A. Simple controller 2	1.34	4.21	6.76
B. Simple controller 1	1.30	2.48	5.55
C. Proposed controller	1.11	1.08	1.24
Improvement (%) B \rightarrow C	14.62	56.45	77.66
Improvement (%) A \rightarrow C	17.16	74.35	81.66

The reserve MAPE of simple controller 1 is already quite low in absolute terms, if the frequency deviation bias is small. Even in this case, however, the improvement achieved by the proposed controller is valuable in order to pass the pre-qualification tests that many system operators apply for reserve providers. In addition, the proposed controller will increase the profits from offering high-quality PFC reserves in a reserve market with performance-based payments.

Figure 9 presents time-domain simulation results to compare the proposed controller against simple controller 1. Specifically, we present the time series of frequency deviation, refrigerator aggregation’s power response, and aggregation’s mean temperature. As shown in Figs. 9b and 9d, there are notable offsets (long-term accumulated errors) in the response of simple controller 1, but not in that of the proposed controller. This significant improvement in “continuous operation” of PFC with refrigerators is due to our contributions in control design. On the other hand, as shown in Fig. 9e, the instantaneous “dynamic response” of the proposed controller is only marginally better than that of simple controller 1. Moreover, the

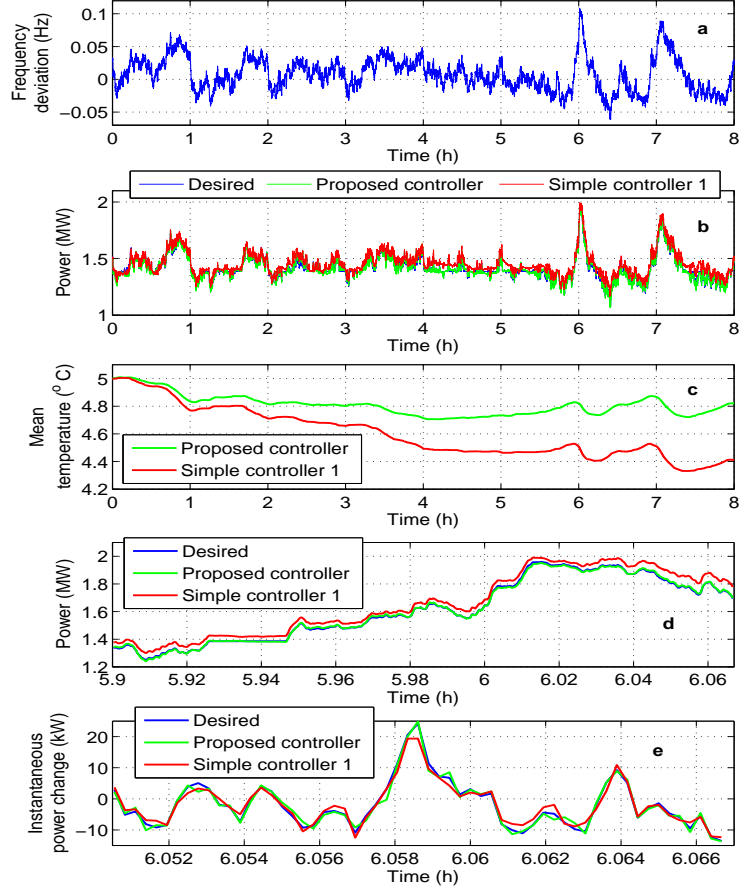


Figure 9: Comparison of the proposed controller against simple controller 1: (a) Frequency deviation; (b) Aggregate refrigerator power; (c) Mean temperature in the aggregation; (d) 10-minute close-up of plot (b); (e) 1-minute close-up of plot (d) with the instantaneous power change to provide PFC reserve.

proposed controller improves user utility by maintaining the aggregation’s mean temperature close to the nominal value of 5°C (see Fig. 9c). This is very important in practice to ensure the participants’ commitment to the demand response program.

8 Robustness Analysis

8.1 Performance with Biased Frequency

In this section we present results for the corrective temperature control loop. Figure 10 compares the reserve error $e_{r,t} = 100 \cdot (P_{d,t} - P_{agg,t})/P_{res}$ and the temperature deviations for three K_c values, as well as for a case without corrective control.




Figure 10: Time-domain simulation results of reserve error and mean temperature deviation for four different values of the corrective temperature gain.




Figure 11: Reserve and temperature errors for 10 corrective temperature gains (averaged over 10 simulations cases).

The results correspond to simulations with a frequency deviation signal that is positively biased for the first 15 hours (with an average bias of $\delta = 0.0192$ Hz) and zero-mean for the rest of the day. If $K_c = 0$, there is a steady-state error in temperature, which has a serious impact on user utility. At the same time, the temperature error affects the baseline power of the population, which results in a large PFC error and reduces the population's ability to provide PFC. On the other hand, all three non-zero gains bring the mean temperature close to the nominal value of 5°C and eliminate the steady-state reserve error. As expected, higher K_c values allow a faster temperature recovery but introduce oscillations in the baseline power, which translate into increased reserve errors at the beginning of the recovery period (hour 5 in Fig. 10).

Results from 10 day-long simulations for K_c values in the range $[0.1, 1] \cdot 10^{-4}$ are presented in Fig. 11. The frequency deviation is positively biased in 4 of the days, negatively biased in another 4 of the days, and zero-mean in the remaining 2 days. The blue bars show the average values of reserve MAPE and temperature




Figure 12: The PDF of temperature deviations at the end of simulation.

Root Mean Square Error (RMSE), whereas the red lines indicate the range spanned from the minimum to the maximum value. As expected, the reserve MAPE is a convex function of K_c . The gain $K_c = 0.5 \cdot 10^{-4}$ results in the minimum reserve MAPE of 1.15%, which is very close to the baseline MAPE of 0.85% due to the system's natural dynamics (shown with the green bar in Fig. 11).

The results of Fig. 10 are in agreement with Proposition 2. With $\delta = 0.0192$ Hz, $N_{ev} = 15$ h, $N_{rec} = 9$ h, $\varepsilon = 1^\circ\text{C}$, $\epsilon = 0.2^\circ\text{C}$, $D^r = 0.15$, $\bar{P}_n = 80$ W and $\bar{\beta} = 4.4 \cdot 10^{-5}$, the lower bound on K_c is $0.4863 \cdot 10^{-4}$. From Fig. 10 we can see that $K_c = 0.5 \cdot 10^{-4}$ is consistent with the design criteria: (i) the maximum temperature deviation is approximately $\varepsilon = 1^\circ\text{C}$; and (ii) $N_{rec} = 9$ h after the frequency bias has disappeared, the mean temperature is less than $\epsilon = 0.2^\circ\text{C}$ away from the nominal value of 5°C . In addition, the results of Fig. 11 are in agreement with (35). For the considered values $\bar{\Delta T} = 2^\circ\text{C}$, $\bar{T}_a = 22^\circ\text{C}$, $\bar{\eta} \bar{R} \bar{P}_n = 70^\circ\text{C}$ and $\Delta t = 1$ s, (35) provides us with the upper bound $K_c = 0.5004 \cdot 10^{-4}$. Therefore, the theoretical lower and upper bounds are very tight.

8.2 Effect of Limited Thermostat Resolution

To investigate the effect of minimum thermostat resolution ΔT_{res} on temperature deviations and control accuracy, we performed 10 day-long simulations with 9 different ΔT_{res} values and present the results in Fig. 12. The plot shows the empirical PDFs of temperature deviations at the end of the simulation, which match closely normal distributions. The mean temperature deviation is independent of ΔT_{res} , whereas the variance increases with ΔT_{res} , i.e., larger deviations are possible for higher ΔT_{res} values. These simulation results are in agreement with Proposition 3 and Remark 2.

Figure 13 compares the following cases: (i) $\Delta T_{res} = 0^\circ\text{C}$, (ii) $\Delta T_{res} = 0.1^\circ\text{C}$ and no bound ΔT_b , and (iii) $\Delta T_b = 0.1^\circ\text{C}$ with $\Delta T_b = 1^\circ\text{C}$. The left plot shows the histogram of temperature limit deviations at the end of a 5-day simulation, whereas the right plot shows the evolution of the standard deviation of temperature

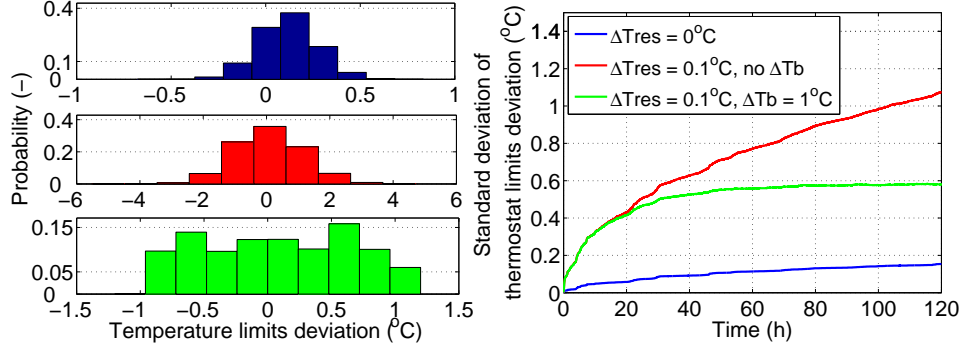


Figure 13: The effect of limited thermostat resolution on the temperature deviations, with and without a temperature deviation bound.




Figure 14: The dependence of reserve MAPE on the thermostat resolution (results averaged over 10 simulations cases).

limit changes over time. If $\Delta T_{res} = 0^\circ\text{C}$ (blue color), the histogram resembles a normal distribution and the standard deviation slowly increases and settles to a small value, as explained in Remark 2. If $\Delta T_{res} = 0.1^\circ\text{C}$ and no ΔT_b is used (red color), the histogram resembles again a normal distribution but the standard deviation monotonically increases, as shown in Proposition 3. However, if a bound $\Delta T_b = 1^\circ\text{C}$ is used (green color), the standard deviation is capped at a significantly lower value. Note that the bounding results in a more uniform PDF of temperature deviations in the population.

The effect of minimum thermostat resolution on control accuracy is shown in Fig. 14. The blue bars are mean values of reserve MAPE over 10 simulations, whereas the red lines indicate the minimum and maximum values. The reserve MAPE generally increases as ΔT_{res} increases but the increase is rather small, i.e., thermostat resolution limitations only slightly reduce control accuracy. In addition, we compared the reserve MAPE for $\Delta T_{res} = 0.1^\circ\text{C}$, with and without a bound $\Delta T_b = 1^\circ\text{C}$, and no observable reduction in control accuracy was found.




Figure 15: Left: The histogram of door openings within a day. Right: Door opening events for a refrigerator generated based on the histogram.

Table 9: Baseline and Reserve MAPE with and without Door Openings

Baseline MAPE (%)		Reserve MAPE (%) Small bias		Reserve MAPE (%) Large bias	
Without	With	Without	With	Without	With
0.90	0.47	1.01	1.41	1.15	2.06

8.3 Performance with Door Openings

In this section, we investigate the effect of refrigerator door openings on controller’s performance. The left plot of Fig. 15 shows the probability of door openings for each hour of the day based on data from [41]. The right plot shows a time series example of door opening events for a refrigerator, which is generated using the probability profile of the left plot.

We performed simulations for the base case scenario with and without door openings and compare the control performance in Table 9. Due to door openings, the reserve MAPE increases by 40% for small frequency deviation bias and by 80% for large frequency deviation bias. However, in both cases the error is still small and therefore the proposed controller is applicable despite the random door openings. The door openings smoothen the baseline power during load ramps, especially around 3 pm and 7 pm, and as a result the baseline MAPE decreases. To illustrate this, Fig. 16 shows the baseline power with and without door openings, as well as the aggregate power without PFC. Notice that there is less natural fluctuation around the baseline with door openings.

Figure 17 shows the PFC activation (“desired power”) around the baseline,, with and without door openings, for a period of 1 hour. The baseline with door openings has an increasing trend, but the PFC activation is very close to that without door openings. Notice that the door openings do not have a significant effect on the tracking of the desired power. Therefore, if the door openings are considered in the control design as shown in the paper, then their effect on the tracking performance of the desired power will not be detrimental.




Figure 16: 24-hour baseline power with and without refrigerator door openings. The minimum baseline power of 1.38 MW corresponds to a duty cycle of 0.25, whereas the maximum power of 2.24 MW to a duty cycle of 0.4.




Figure 17: PFC activation around the baseline without refrigerator door openings (left) and with refrigerator door openings (right) for a period of 1 hour.

9 Implementation Issues

9.1 Refrigerator Model

Our analysis is based on the commonly used first-order freezer-less refrigerator model (1), (2). As shown in [40], using separate states for the temperature of the refrigerator compartment's air, refrigerator's content and evaporator, significantly improves model performance. Modeling a refrigerator with a freezer is possible following the approach of [10].

It is possible to integrate more detailed refrigerator models in the control design by modifying (29), (30) and (31). More importantly, the model of startup dynamics can be improved by considering the dependence on evaporator temperature when a refrigerator switches on. Since this temperature depends on the time elapsed since the previous on cycle, which is not constant when providing PFC, the profile of

startup dynamics will likely be time-varying in reality.

9.2 Noisy Frequency Measurements

The authors of [44] performed a correlation analysis of frequency measurements at the transmission level (230 kV) and at the wall outlet level (120 V). The results were promising for PFC and showed that the frequency measurements at the two voltage levels are highly correlated. Although it is important to investigate how measurement noise affects control performance, a moving average filter can be used to reduce the measurement noise in a practical implementation [23].

9.3 Cost and Economic Evaluation

If $D^r = 0.2$, our simulations show that 63,500 refrigerators are necessary to provide 1 MW of PFC reserves.⁶ Using a reserve capacity price of 21.5 e /MW/h (average price from January to September 2014 in Switzerland), the total revenue from selling 1 MW of capacity is 188,340 e /year, which corresponds to a revenue per device as low as 3 e /year. Since the average refrigerator lifetime is 14 years, providing PFC reserves will be profitable only if the implementation cost is below 42 e .

We expect the micro-controller's cost to be a few e due to its simplicity, whereas the frequency meter is the dominant cost driver of the proposed scheme. Reference [23] reports a cost of 20 e for a controller including the frequency meter, but without mentioning the meter's accuracy. According to our market search, a typical selling price in 2016 for a frequency meter with accuracy 10 mHz was approximately 40 e , whereas the price explodes to approximately 600 e for an accuracy of 1 mHz. Based on discussions with manufacturers of high-accuracy frequency meters, the cost is expected to drop to approximately 50 e , if the meters are standardized and integrated in the refrigerators' electronic circuits.

Nevertheless, even with a reduced frequency meter cost, providing PFC with residential refrigerators will not be a very profitable investment due to the required large number of devices and the currently low reserve prices. In order to make customer participation in such programs more attractive, the refrigerator owners could enjoy electricity price rebates. In addition, performance-based reserve payments could be used in order to remunerate the refrigerators' higher accuracy and faster response in reserve provision. In fact, performance-based payments are already used for secondary frequency reserves in the PJM power market in the US.

The developed methods are also applicable to other cooling devices, for example, commercial refrigerators and freezers, refrigerated warehouses and air conditioners, with no or minimal modifications. Due to the much higher power rating of these devices, a significantly smaller aggregation will be needed for the same

⁶Assuming a refrigerator power rating of 80 W and a duty cycle of 0.3, the average power consumption of the aggregation of 63,500 devices is 1.524 MW. Therefore, the aggregation provides $100 \cdot (1/1.524) \approx 65\%$ of the average power as reserve capacity.

amount of reserves. Therefore, we expect the revenue for each participant to be notably higher, which makes this application case financially more attractive compared with residential refrigerators.

10 Concluding Remarks

In this paper, we presented a decentralized stochastic control scheme to enable PFC reserve provision from aggregations of refrigerators. The control is based on probabilistic switching to avoid load synchronization, and it accounts for refrigerator startup dynamics and lockout constraints. Reliable reserve provision during persistent frequency deviations is achieved by resetting the thermostat temperature limits. Furthermore, the controller is robust to biased frequency deviations, limited thermostat resolution and door openings. Extensive simulation results show that the control scheme allows a refrigerator aggregation to provide fast and reliable PFC without communication. In the future, we plan to investigate the controller's robustness to excessive compressor locking, and perform dynamic simulation studies in a two-area power system model.

Appendix: Proof of Proposition 1

Proof. The estimated aggregate power at time step t after reserve activation is

$$\hat{P}_{\text{agg},t} = N_{\text{r,a}} \bar{P}_{\text{n}} \bar{S}_{\text{u}}(t), \quad (54)$$

where $N_{\text{r,a}}$ is the number of activated devices. Assuming independence of random variables $P_{\text{n},i}$, u_i and $N_{\text{s},i}$, the expected actual aggregate power is computed as

$$\begin{aligned} \bar{P}_{\text{agg},t} &= \mathbb{E} \left[\sum_{i=1}^{N_{\text{r,a}}} P_{i,t} \right] = \sum_{i=1}^{N_{\text{r,a}}} \mathbb{E} [P_{i,t}] \\ &= \sum_{i=1}^{N_{\text{r,a}}} \mathbb{E} \left[P_{\text{n},i} \left[1 + u_i \left[1 - \frac{t}{N_{\text{s},i}} \right]_+ \right] \right] \\ &= \sum_{i=1}^{N_{\text{r,a}}} \bar{P}_{\text{n}} \left[1 + \bar{u} \cdot \mathbb{E} \left[1 - \frac{t}{N_{\text{s},i}} \right]_+ \right]. \end{aligned} \quad (55)$$

Part 1: If $t \leq N_{\text{s,min}}$, then $[\cdot]_+$ is redundant and (55) gives

$$\begin{aligned} \bar{P}_{\text{agg},t} &= \sum_{i=1}^{N_{\text{r,a}}} \bar{P}_{\text{n}} \left[1 + \bar{u} \left[1 - \mathbb{E} \left[\frac{t}{N_{\text{s},i}} \right] \right] \right] \\ &\leq \sum_{i=1}^{N_{\text{r,a}}} \bar{P}_{\text{n}} \left[1 + \bar{u} \left(1 - \frac{t}{\bar{N}_{\text{s}}} \right) \right] = N_{\text{r,a}} \bar{P}_{\text{n}} \bar{S}_{\text{u}}(t) = \hat{P}_{\text{agg},t}, \end{aligned} \quad (56)$$

where Jensen's inequality $\phi(\mathbb{E}[N_{s,i}]) \leq \mathbb{E}[\phi(N_{s,i})]$ is invoked with $\phi(N_{s,i}) = t/N_{s,i}$, which is a convex function since $t, N_{s,i} \geq 0$. Thus, $\bar{P}_{\text{agg},t} \leq \hat{P}_{\text{agg},t}$ holds for $t \leq N_{s,\min}$.

Part 2: For $t \geq \bar{N}_s$, we obtain $\hat{P}_{\text{agg},t} = 0$ from (9) and (54). Since the expected aggregate power $\bar{P}_{\text{agg},t}$ is non-negative, $\hat{P}_{\text{agg},t}$ is clearly a lower bound.

If $N_{s,i}$ follows the uniform PDF $N_{s,i} \sim [N_{s,\min}, N_{s,\max}]$, then for a given $t \in (N_{s,\min}, \bar{N}_s)$ the probability that a randomly selected refrigerator contributes to the aggregate startup dynamics is equal to $\mathbb{P}(N_{s,i} \geq t) = (N_{s,\max} - t) / (N_{s,\max} - N_{s,\min})$. Let us denote by $\tilde{N}_{s,i}$ the random startup duration that follows the uniform PDF $\tilde{N}_{s,i} \sim [t, N_{s,\max}]$. The expectation from the last step of (55) can be written as

$$\mathbb{E} \left[1 - \frac{t}{N_{s,i}} \right]_+ = \mathbb{E} \left[1 - \frac{t}{\tilde{N}_{s,i}} \right] \cdot \frac{N_{s,\max} - t}{N_{s,\max} - N_{s,\min}} \quad (57a)$$

$$\leq \left(1 - \frac{t}{\mathbb{E}[\tilde{N}_{s,i}]} \right) \cdot \frac{N_{s,\max} - t}{N_{s,\max} - N_{s,\min}} \quad (57b)$$

$$= \left[1 - \frac{t}{0.5 \cdot (t + N_{s,\max})} \right] \cdot \frac{N_{s,\max} - t}{N_{s,\max} - N_{s,\min}} \quad (57c)$$

$$= \frac{(N_{s,\max} - t)^2}{(N_{s,\max} - N_{s,\min}) \cdot (N_{s,\max} + t)}. \quad (57d)$$

We used Jensen's inequality in (57b) and the fact that $\tilde{N}_{s,i}$ follows the uniform distribution $\tilde{N}_{s,i} \sim [t, N_{s,\max}]$ in step (57c).

The inequality $\hat{P}_{\text{agg},t} \geq \bar{P}_{\text{agg},t}$ holds if and only if

$$1 - \frac{t}{\bar{N}_s} \geq \mathbb{E} \left[1 - \frac{t}{N_{s,i}} \right]_+ \Leftrightarrow \quad (58)$$

$$\frac{N_{s,\max} + N_{s,\min} - 2t}{N_{s,\max} + N_{s,\min}} \geq \frac{(N_{s,\max} - t)^2}{(N_{s,\max} - N_{s,\min}) \cdot (N_{s,\max} + t)}, \quad (59)$$

where we used the fact that $\bar{N}_s = 0.5 \cdot (N_{s,\max} + N_{s,\min})$. With some effort (59) can be equivalently written as

$$(N_{s,\min} - 3N_{s,\max}) \cdot t^2 + (N_{s,\max}^2 + 4N_{s,\min}N_{s,\max} - N_{s,\min}^2) \cdot t - (N_{s,\min}^2 N_{s,\max} + N_{s,\min} N_{s,\max}^2) \geq 0. \quad (60)$$

The quadratic form in (60) has two real solutions t_1 and t_{lim} . The first one is $t_1 = N_{s,\min}$ by inspection. For the second one, we use Vieta's formula $t_1 \cdot t_{\text{lim}} = c/a$ with $a = N_{s,\min} - 3N_{s,\max}$ and $c = -(N_{s,\min}^2 N_{s,\max} + N_{s,\min} N_{s,\max}^2)$ and find

$$t_{\text{lim}} = \frac{N_{s,\max} (N_{s,\min} + N_{s,\max})}{3N_{s,\max} - N_{s,\min}}. \quad (61)$$

The quadratic inequality (60) holds for $t \in [t_1, t_{\text{lim}}]$. Therefore, if $t > N_{s,\min}$ the inequality $\hat{P}_{\text{agg},t} \geq \bar{P}_{\text{agg},t}$ holds for $t \leq t_{\text{lim}}$. ■

Appendix: Proof of Proposition 2

Proof. Part 1: For $0 \leq t \leq N_{\text{ev}} + 1$ using $\bar{T}_0 = \bar{T}_{\text{nom}}$ and (38), equation (37) gives

$$\bar{T}_t = \bar{T}_{\text{nom}} - \gamma\delta \cdot \sum_{k=0}^{t-1} \lambda^k = \bar{T}_{\text{nom}} - \gamma\delta \cdot \frac{1 - \lambda^{t-1}}{1 - \lambda}, \quad (62)$$

where a standard property of geometric series is invoked. The minimum of (62) is obtained for $t = N_{\text{ev}} + 1$ and thus if $\bar{T}_{\text{nom}} - \varepsilon \leq \bar{T}_t \leq \bar{T}_{\text{nom}} + \varepsilon$ holds for $t = N_{\text{ev}} + 1$ then it holds $\forall t$. Assuming $\delta > 0$ without loss of generality, it should hold

$$\bar{T}_{t=N_{\text{ev}}+1} \geq \bar{T}_{\text{nom}} - \varepsilon \Rightarrow -\gamma\delta \cdot \frac{1 - \lambda^{N_{\text{ev}}}}{1 - \lambda} \geq -\varepsilon. \quad (63)$$

Part 2: For $t \geq N_{\text{ev}} + 1$, $\Delta f_t = 0$ and the mean temperature recovers towards \bar{T}_{nom} . We reset the time index such that $t' = 0 \Leftrightarrow t = N_{\text{ev}} + 1$. The new initial condition is $\bar{T}_{(t'=0)} = \bar{T}_{(t=N_{\text{ev}}+1)}$ and is computed from (62). Equation (37) gives

$$\bar{T}_{t'} = \bar{T}_{\text{nom}} - \gamma\delta \cdot \frac{1 - \lambda^{N_{\text{ev}}}}{1 - \lambda} \cdot \lambda^{t'}. \quad (64)$$

At time $t' = N_{\text{rec}}$, $\bar{T}_{\text{nom}} - \epsilon \leq \bar{T}_{N_{\text{rec}}} \leq \bar{T}_{\text{nom}} + \epsilon$ must hold. Assuming $\delta > 0$ without loss of generality, we get

$$\bar{T}_{t'=N_{\text{rec}}} \geq \bar{T}_{\text{nom}} - \epsilon \Rightarrow -\gamma\delta \cdot \lambda^{N_{\text{rec}}} \cdot \frac{1 - \lambda^{N_{\text{ev}}}}{1 - \lambda} \geq -\epsilon. \quad (65)$$

Equations (63) and (65) are equivalent to (39) and (40). ■

Appendix: Proof of Proposition 3

Proof. The change in temperature limits of a randomly selected refrigerator i can be expressed as

$$\Delta T_{\text{lim},i,t} = \sum_{k=0}^{t-1} \Delta T_{\text{res}} \cdot \zeta_k \cdot z_k, \quad (66)$$

where ζ_k is equal to 1 or -1 depending on the sign of $\Delta T_{\text{lim},k}$, and z_k is the Bernoulli random variable

$$f_z(z_k) = \begin{cases} \varrho_k, & \text{if } z_k = 1 \\ 1 - \varrho_k, & \text{if } z_k = 0 \end{cases}. \quad (67)$$

Using the fact $\mathbb{E}[z_k] = \varrho_k$, the mean value is computed as

$$\begin{aligned} \mathbb{E}[\Delta T_{\text{lim},i,t}] &= \sum_{k=0}^{t-1} \Delta T_{\text{res}} \zeta_k \mathbb{E}[z_k] = \sum_{k=0}^{t-1} \Delta T_{\text{res}} \zeta_k \varrho_k \\ &= \sum_{k=0}^{t-1} \zeta_k |\Delta T_{\text{lim},k}| = \sum_{k=0}^{t-1} \Delta T_{\text{lim},k}. \end{aligned} \quad (68)$$

The variance is computed as

$$\text{Var} [\Delta T_{\text{lim},i,t}] = \text{Var} \left[\sum_{k=0}^{t-1} \Delta T_{\text{res}} \cdot \zeta_k \cdot z_k \right] \quad (69a)$$

$$= \sum_{k=0}^{t-1} \Delta T_{\text{res}}^2 \cdot \zeta_k^2 \cdot \text{Var} [z_k] = \sum_{k=0}^{t-1} \Delta T_{\text{res}}^2 \cdot \varrho_k \cdot (1 - \varrho_k) \quad (69b)$$

$$= \sum_{k=0}^{t-1} \Delta T_{\text{res}}^2 \cdot \frac{|\Delta T_{\text{lim},k}|}{\Delta T_{\text{res}}} \cdot \left(1 - \frac{|\Delta T_{\text{lim},k}|}{\Delta T_{\text{res}}} \right) \quad (69c)$$

$$= \sum_{k=0}^{t-1} |\Delta T_{\text{lim},k}| \cdot (\Delta T_{\text{res}} - |\Delta T_{\text{lim},k}|), \quad (69d)$$

where the Bernoulli random variables z_k and z_j are assumed to be uncorrelated $\forall k \neq j$ in (69a), $\text{Var} [z_k] = \varrho_k \cdot (1 - \varrho_k)$ is used in (69b), and the definition of ϱ_k from (42) is used in (69c). ■

Appendix: Proof of Proposition 4

Proof. Denote by E_{op}^{d} and E_{cl}^{d} the refrigerator's daily energy consumption with and without door openings, respectively. Consider a refrigerator subject to door openings and denote by \bar{T}_{op} and \bar{T}_{cl} the time-averaged temperature while the door is open and closed, respectively. The refrigerator operates on average $\mu_{\text{op}} \cdot \mu_{\text{d}}$ seconds with open door, and $N^{\text{d}} - \mu_{\text{op}} \cdot \mu_{\text{d}}$ seconds with closed door. The refrigerator's time constant with closed door is $\alpha = 1/(RC)$, whereas with open door it is $\alpha_{\text{op}} = 1/(R_{\text{op}}C)$.

The refrigerator's energy consumption is equal to the integral of instantaneous thermal losses over time because its temperature is regulated. Assuming that \bar{T}_{cl} and \bar{T}_{op} are constant, the daily energy consumption with and without door openings can be approximated using (1), (2) as

$$E_{\text{cl}}^{\text{d}} \approx (\alpha/\beta) \cdot (T_{\text{a}} - \bar{T}_{\text{cl}}) \cdot N^{\text{d}} \quad (70)$$

$$E_{\text{op}}^{\text{d}} \approx (\alpha/\beta) \cdot (T_{\text{a}} - \bar{T}_{\text{cl}}) \cdot (N^{\text{d}} - \mu_{\text{op}}\mu_{\text{d}}) + \quad (71)$$

$$(\alpha_{\text{op}}/\beta) \cdot (T_{\text{a}} - \bar{T}_{\text{op}}) \cdot \mu_{\text{op}}\mu_{\text{d}}, \quad (72)$$

whereas E_{op}^{d} and E_{cl}^{d} are related according to

$$E_{\text{op}}^{\text{d}} = (1 + \xi) \cdot E_{\text{cl}}^{\text{d}}. \quad (73)$$

Combining (70), (72) and (73), we get

$$\begin{aligned} (\alpha/\beta)(T_{\text{a}} - \bar{T}_{\text{cl}})(N^{\text{d}} - \mu_{\text{op}}\mu_{\text{d}}) + (\alpha_{\text{op}}/\beta)(T_{\text{a}} - \bar{T}_{\text{op}})\mu_{\text{op}}\mu_{\text{d}} = \\ (1 + \xi) \cdot (\alpha/\beta) \cdot (T_{\text{a}} - \bar{T}_{\text{cl}}) \cdot N^{\text{d}}, \end{aligned} \quad (74)$$

which when solved for α_{op} gives

$$\alpha_{\text{op}} = \alpha \cdot \frac{T_{\text{a}} - \bar{T}_{\text{cl}}}{T_{\text{a}} - \bar{T}_{\text{op}}} \cdot \left(1 + \frac{N^{\text{d}}}{\mu_{\text{op}} \cdot \mu_{\text{d}}} \cdot \xi \right) \Rightarrow \quad (75)$$

$$\alpha_{\text{op}} \geq \alpha \cdot [1 + (N^{\text{d}} \cdot \xi)/(\mu_{\text{op}} \cdot \mu_{\text{d}})], \quad (76)$$

because $\bar{T}_{cl} \leq \bar{T}_{op}$. Inequality (46) now directly follows from (76) using $\alpha = (1/RC)$ and $\alpha_{op} = (1/R_{op}C)$. ■

References

- [1] Y. Makarov, C. Loutan, J. Ma, and P. De Mello, “Operational impacts of wind generation on California power systems,” *IEEE Transactions on Power Systems*, vol. 24, no. 2, pp. 1039–1050, 2009.
- [2] D. Callaway and I. Hiskens, “Achieving controllability of electric loads,” *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, 2011.
- [3] F. Schweppe, R. Tabors, J. Kirtley Jr, H. Outhred, F. Pickel, and A. Cox, “Homeostatic utility control,” *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-99, no. 3, pp. 1151–1163, 1980.
- [4] C. Chong and A. Debs, “Statistical synthesis of power system functional load models,” in *18th IEEE Conference on Decision and Control including the Symposium on Adaptive Processes*, no. 18, 1979, pp. 264–269.
- [5] R. Malhame and C.-Y. Chong, “Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system,” *IEEE Transactions on Automatic Control*, vol. 30, no. 9, pp. 854–860, 1985.
- [6] F. Constantopoulos, P. Schweppe and R. Larson, “Estia: A real-time consumer control scheme for space conditioning usage under spot electricity pricing,” *Computers and Operations Research*, vol. 18, no. 8, pp. 751–765, Jan. 1991.
- [7] D. S. Callaway, “Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy,” *Energy Conversion and Management*, vol. 50, no. 5, pp. 1389–1400, 2009.
- [8] C. Perfumo, E. Kofman, J. H. Braslavsky, and J. K. Ward, “Load management: Model-based control of aggregate power for populations of thermostatically controlled loads,” *Energy Conversion and Management*, vol. 55, pp. 36–48, 2012.
- [9] S. Kundu, N. Sinitsyn, S. Backhaus, and I. Hiskens, “Modeling and control of thermostatically controlled loads,” *arXiv preprint arXiv:1101.2157*, 2011.
- [10] T. M. Keep, F. E. Sifuentes, D. M. Auslander, and D. S. Callaway, “Using load switches to control aggregated electricity demand for load following and regulation,” in *IEEE PES General Meeting (GM)*, Detroit, USA, July 2011.
- [11] W. Zhang, J. Lian, C.-Y. Chang, and K. Kalsi, “Aggregated modeling and control of air conditioning loads for demand response,” *IEEE Transactions on Power Systems*, vol. 28, no. 4, pp. 4655–4664, Nov 2013.

- [12] L. Totu and R. Wisniewski, "Demand response of thermostatic loads by optimized switching-fraction broadcast," in *19th IFAC World Congress*, Cape Town, South Africa, August 2014.
- [13] C. Moya, W. Zhang, J. Lian, and K. Kalsi, "A hierarchical framework for demand-side frequency control," in *American Control Conference (ACC)*, Portland, USA, June 2014, pp. 52–57.
- [14] S. Koch, J. Mathieu, and D. Callaway, "Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services," in *Proceedings of the Power Systems Computation Conference*, Stockholm, Sweden, Aug. 2011.
- [15] J. Kondoh, N. Lu, and D. Hammerstrom, "An evaluation of the water heater load potential for providing regulation service," *IEEE Transactions on Power Systems*, vol. 26, no. 3, pp. 1309–1316, 2011.
- [16] E. Vrettos, S. Koch, and G. Andersson, "Load frequency control by aggregations of thermally stratified electric water heaters," in *IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe)*, Berlin, Germany, Oct. 2012.
- [17] J. Mathieu, S. Koch, and D. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Transactions on Power Systems*, vol. 28, no. 1, pp. 430–440, 2013.
- [18] E. Can Kara, Z. Kolter, M. Berges, B. Krogh, G. Hug, and T. Yuksel, "A moving horizon state estimator in the control of thermostatically controlled loads for demand response," in *Proceedings of the IEEE SmartGridComm Conference*, Vancouver, Canada, October 2013.
- [19] T. Borsche, F. Oldewurtel, and G. Andersson, "Minimizing communication cost for demand response using state estimation," in *Proceedings of the PowerTech Conference*, Grenoble, France, June 2013.
- [20] S. Moura, J. Bendtsen, and V. Ruiz, "Observer design for boundary coupled PDEs: Application to thermostatically controlled loads in smart grids," in *Proceedings of the IEEE Conference on Decision and Control*, Florence, Italy, Dec. 2013.
- [21] E. Vrettos, J. L. Mathieu, and G. Andersson, "Control of thermostatic loads using moving horizon estimation of individual load states," in *Power Systems Computation Conference (PSCC)*, Wroclaw, Poland, Aug. 2014.
- [22] J. Short, D. Infield, and L. Freris, "Stabilization of grid frequency through dynamic demand control," *IEEE Transactions on Power Systems*, vol. 22, no. 3, pp. 1284–1293, Aug. 2007.

- [23] Z. Xu, J. Østergaard, and M. Tøgeby, "Demand as frequency controlled reserve," *IEEE Transactions on Power Systems*, vol. 26, no. 3, pp. 1062–1071, 2011.
- [24] A. Molina-Garcia, F. Bouffard, and D. S. Kirschen, "Decentralized demand-side contribution to primary frequency control," *IEEE Transactions on Power Systems*, vol. 26, no. 1, pp. 411–419, 2011.
- [25] K. Samarakoon, J. Ekanayake, and N. Jenkins, "Investigation of domestic load control to provide primary frequency response using smart meters," *IEEE Transactions on Smart Grid*, vol. 3, no. 1, pp. 282–292, 2012.
- [26] D. Angeli and P. A. Kountouriotis, "A stochastic approach to dynamic-demand refrigerator control," *IEEE Transactions on Control Systems Technology*, vol. 20, no. 3, pp. 581–592, May 2012.
- [27] E. Kremers, J. Mari, and O. Barambones, "Emergent synchronisation properties of a refrigerator demand side management system," *Applied Energy*, vol. 101, pp. 709–717, 2013.
- [28] D. Hammerstrom *et al.*, "Pacific northwest gridwise testbed demonstration projects - Part II: Grid friendly appliance project," Pacific Northwest National Laboratory, Tech. Rep., 2007.
- [29] V. Moghadam, R. Mohammad, R. T. Ma, and R. Zhang, "Distributed frequency control in smart grids via randomized demand response," *IEEE Transactions on Smart Grid*, vol. 5, no. 6, pp. 2798–2809, 2014.
- [30] S. H. Tindemans, V. Trovato, and G. Strbac, "Decentralized control of thermostatic loads for flexible demand response," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 5, pp. 1685–1700, 2015.
- [31] T. Borsche, U. Markovic, and G. Andersson, "A new algorithm for primary frequency control with cooling appliances," *Computer Science-Research and Development*, pp. 1–7, 2015.
- [32] B. Biegel, L. H. Hansen, P. Andersen, and J. Stoustrup, "Primary control by on/off demand-side devices," *IEEE Transactions on Smart Grid*, vol. 4, no. 4, pp. 2061–2071, 2013.
- [33] S. Weckx, J. Driesen, and R. D'Hulst, "Optimal frequency support by dynamic demand," in *PowerTech Conference*, Grenoble, France, June 2013, pp. 1–6.
- [34] C. Zhao, U. Topcu, and S. H. Low, "Frequency-based load control in power systems," in *American Control Conference (ACC)*, 2012, Montreal, Canada, June 2012, pp. 4423–4430.

- [35] C. Zhao, U. Topcu, N. Li, and S. Low, “Design and stability of load-side primary frequency control in power systems,” *IEEE Transactions on Automatic Control*, vol. 59, no. 5, pp. 1177–1189, 2014.
- [36] C. Ziras, E. Vrettos, and G. Andersson, “Primary frequency control with refrigerators under startup dynamics and lockout constraints,” in *IEEE PES General Meeting*, Denver, USA, July 2015.
- [37] J. M. Goncalves, C. Melo, C. J. L. Hermes, and J. R. Barbosa Jr., “Experimental mapping of the thermodynamic losses in vapor compression refrigeration systems,” *Journal of the Brazilian Society of Mechanical Sciences and Engineering*, vol. 33, pp. 159 – 165, June 2011.
- [38] H. K. Khalil and J. Grizzle, *Nonlinear systems*. Prentice hall New Jersey, 1996, vol. 3.
- [39] G. Van Kempen and L. Van Vliet, “Mean and variance of ratio estimators used in fluorescence ratio imaging,” *Cytometry*, vol. 39, no. 4, pp. 300–305, 2000.
- [40] F. Sossan, V. Lakshmanan, G. T. Costanzo, M. Marinelli, P. J. Douglass, and H. Bindner, “Grey-box modelling of a household refrigeration unit using time series data in application to demand side management,” *Sustainable Energy, Grids and Networks*, vol. 5, pp. 1–12, 2016.
- [41] D. Parker and T. Stedman, “Measured electricity savings of refrigerator replacement: Case study and analysis,” *Proceeding of the 1992 Summer Study on Energy Efficiency in Buildings, American Council for an Energy Efficient Economy*, vol. 3, pp. 199–211, 1992.
- [42] E. Vrettos, C. Ziras, and G. Andersson, “Integrating large shares of heterogeneous thermal loads in power system frequency control,” in *IEEE Powertech Conference*, Eindhoven, the Netherlands, June 2015.
- [43] “Test for primary control capability,” Available online at www.swissgrid.ch.
- [44] M. Donnelly, D. Harvey, R. Munson, and D. Trudnowski, “Frequency and stability control using decentralized intelligent loads: Benefits and pitfalls,” in *IEEE Power and Energy Society GM*, 2010, pp. 1–6.